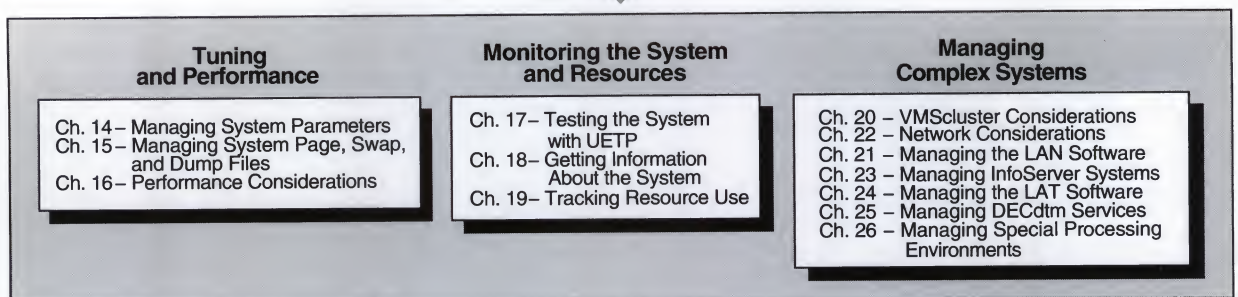
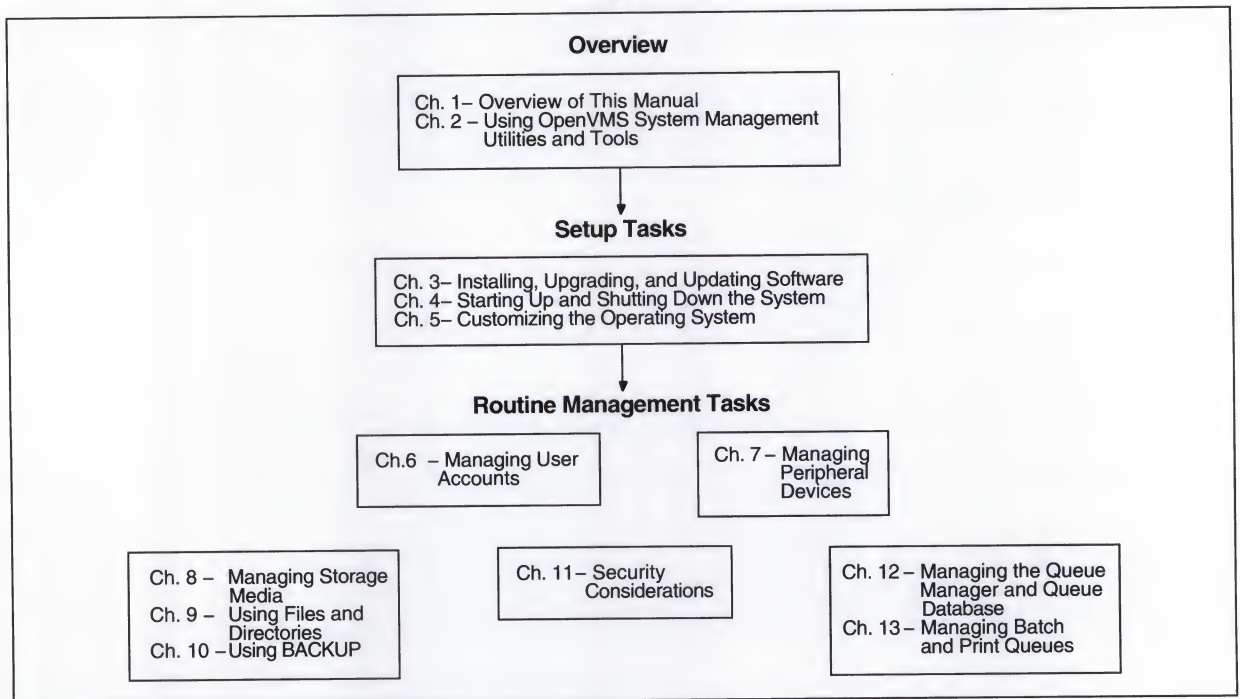



digital

**OpenVMS System Manager's
Manual: Tuning, Monitoring,
and Complex Systems**

OpenVMS





OpenVMS System Manager's Manual: Tuning, Monitoring, and Complex Systems

Order Number: AA-PV5NC-TK

May 1995

This manual is part two of a task-oriented guide to managing an OpenVMS system.

Revision/Update Information: This manual supersedes the *OpenVMS System Manager's Manual: Tuning, Monitoring, and Complex Systems*, OpenVMS AXP Version 6.1 and OpenVMS VAX Version 6.1.

Software Version: OpenVMS Alpha Version 6.2
OpenVMS VAX Version 6.2

**Digital Equipment Corporation
Maynard, Massachusetts**

May 1995

Digital Equipment Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

Possession, use, or copying of the software described in this publication is authorized only pursuant to a valid written license from Digital or an authorized sublicensor.

Digital conducts its business in a manner that conserves the environment and protects the safety and health of its employees, customers, and the community.

© Digital Equipment Corporation May 1995. All rights reserved.

The following are trademarks of Digital Equipment Corporation: ACMS, ALL-IN-1, AXP, Bookreader, CI, DBMS, DDCMP, DECams, DECdtm, DECevent, DECintact, DECnet, DECperformance, DEC RALLY, DECscheduler, DECTalk, DECThreads, DECtrace, DECwindows, DEQNA, Digital, HSC, IAS, InfoServer, LAT, MicroVAX, MicroVAX II, MSCP, OpenVMS, OpenVMS RMS, PATHWORKS, POLYCENTER, PrintServer, Q-bus, Rdb/VMS, RSX-11M, RSX-11M-PLUS, RSX-11D, RT-11, TMSCP, TURBOchannel, VAX, VAXcluster, VAX DOCUMENT, VMS, VMScluster, the Alpha logo, and the DIGITAL logo.

The following are third-party trademarks:

Intel is a trademark of Intel Corporation.

Internet is a registered trademark of Internet, Inc.

ITC Avant Garde Gothic is a registered trademark of International Typeface Corporation.

Microsoft, MS, and MS-DOS are registered trademarks of Microsoft Corporation.

Motif, OSF, OSF/1, OSF/Motif, and Open Software Foundation are registered trademarks of the Open Software Foundation, Inc.

POSIX is a registered trademark of IEEE.

POSTSCRIPT is a registered trademark of Adobe Systems Incorporated.

UNIX is a registered trademark licensed exclusively by X/Open Company Ltd.

Xerox is a registered trademark of Xerox Corporation.

All other trademarks and registered trademarks are the property of their respective holders.

ZK6016

This document is available on CD-ROM.

This document was prepared using VAX DOCUMENT Version 2.1.

Send Us Your Comments

We welcome your comments on this or any other OpenVMS manual. If you have suggestions for improving a particular section or find any errors, please indicate the title, order number, chapter, section, and page number (if available). We also welcome more general comments. Your input is valuable in improving future releases of our documentation.

You can send comments to us in the following ways:

- Internet electronic mail: `openvmsdoc@zko.mts.dec.com`
- Fax: 603-881-0120 Attn: OpenVMS Documentation, ZK03-4/U08
- Online form

Print or edit the online form `SY$HELP:OPENVMSDOC_COMMENTS.TXT`. Send the completed online form by electronic mail to our Internet address, or send the completed hardcopy form by fax or through the postal service.

Please send letters or the form to:

Digital Equipment Corporation
Information Design and Consulting
OpenVMS Documentation
110 Spit Brook Road, ZK03-4/U08
Nashua, NH 03062-2698
USA

Thank you.

How To Order Additional Documentation

Use the following table to order additional documentation or information.
If you need help deciding which documentation best meets your needs, call
800-DIGITAL (800-344-4825).

Telephone and Direct Mail Orders

Location	Call	Fax	Write
U.S.A.	DECdirect 800.DIGITAL 800.344.4825	Fax: 800.234.2298	Digital Equipment Corporation P.O. Box CS2008 Nashua, NH 03061
Puerto Rico	809.781.0505	Fax: 809.749.8300	Digital Equipment Caribbean, Inc. 3 Digital Plaza, 1st Street, Suite 200 P.O. Box 11038 Metro Office Park San Juan, Puerto Rico 00910-2138
Canada	800.267.6215	Fax: 613.592.1946	Digital Equipment of Canada, Ltd. Box 13000 100 Herzberg Road Kanata, Ontario, Canada K2K 2A6 Attn: DECdirect Sales
International	—	—	Local Digital subsidiary or approved distributor
Internal Orders	DTN: 264.3030 603.884.3030	Fax: 603.884.3960	U.S. Software Supply Business Digital Equipment Corporation 10 Cotton Road Nashua, NH 03063-1260

ZK-7654A-GE

Contents

Preface	xix
14 Managing System Parameters	
14.1 Understanding System Parameters	14-1
14.1.1 Default, Current, and Active Values	14-3
14.1.2 Pages and Pagelets	14-4
14.2 Recommended Method for Changing Parameter Values	14-4
14.3 Converting Your Customized Parameter Settings for Use with AUTOGEN	14-5
14.4 Understanding the AUTOGEN Command Procedure	14-8
14.4.1 AUTOGEN Feedback	14-10
14.4.2 Feedback Report (AGEN\$PARAMS.REPORT)	14-11
14.4.3 AUTOGEN Phases	14-16
14.4.4 AUTOGEN Parameter File (MODPARAMS.DAT)	14-17
14.5 Modifying System Parameters with AUTOGEN	14-18
14.5.1 Controlling AUTOGEN's Parameter Settings with MODPARAMS.DAT	14-18
14.5.1.1 Increasing a Value with the ADD_ Prefix	14-19
14.5.1.2 Specifying a Minimum Value with the MIN_ Prefix	14-20
14.5.1.3 Specifying a Maximum Value with the MAX_ Prefix	14-20
14.5.1.4 Specifying an Absolute Value	14-20
14.5.1.5 Defining the Number of VAXcluster Nodes (VAX Only)	14-21
14.5.1.6 Defining the Number of Ethernet Adapters (VAX Only)	14-21
14.5.1.7 Presetting Parameter Values Before Adding Memory (VAX Only)	14-21
14.5.1.8 Overriding Parameters Related to DECnet	14-21
14.5.2 Specifying a Minimum Required Age for Feedback (VAX Only)	14-21
14.5.3 Including an External Parameter File in MODPARAMS.DAT	14-22
14.6 Automating AUTOGEN Reports	14-22
14.6.1 Changing Parameter Values After Reviewing AUTOGEN Reports ...	14-25
14.7 Managing System Parameters with the System Management Utility (SYSMAN)	14-25
14.7.1 Understanding Parameter Values and SYSMAN	14-26
14.7.2 Showing Parameter Values with SYSMAN	14-27
14.7.3 Modifying a Parameter File with SYSMAN	14-28
14.7.4 Modifying Active Values with SYSMAN	14-28
14.8 Managing System Parameters with the System Generation Utility (SYSGEN)	14-30
14.8.1 Understanding Parameter Values and SYSGEN	14-30
14.8.2 Showing Parameter Values with SYSGEN	14-31
14.8.3 Modifying the System Parameter File with SYSGEN	14-33
14.8.4 Modifying Active Values with SYSGEN	14-33
14.8.5 Creating a New Parameter File with SYSGEN	14-34

14.9	Modifying System Parameters with a Conversational Boot	14-36
15	Managing System Page, Swap, and Dump Files	
15.1	Understanding the System Dump File	15-2
15.2	Understanding Page and Swap Files	15-4
15.3	Displaying Information About Page and Swap Files	15-6
15.4	Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files	15-6
15.4.1	Calculating System Dump File Size	15-6
15.4.2	Calculating Page File Size	15-8
15.4.3	Calculating Swap File Size	15-9
15.5	Minimizing Dump File Size When Disk Space Is Insufficient	15-10
15.6	Using SDA to Analyze the Contents of a Crash Dump	15-11
15.7	Using SDA CLUE Commands to Analyze Crash Dump Files (Alpha Only)	15-11
15.7.1	Understanding CLUE (Alpha Only)	15-11
15.7.2	Displaying Data Using SDA CLUE Commands (Alpha Only)	15-12
15.8	Using CLUE to Obtain Historical Information About Crash Dumps (VAX Only)	15-12
15.8.1	Understanding CLUE (VAX Only)	15-12
15.8.2	Displaying Data Using CLUE (VAX Only)	15-13
15.9	Copying Dump Files to Tape or Disk	15-14
15.10	Dump File Off the System Disk (VAX Systems Only)	15-14
15.10.1	Requirements	15-14
15.11	Saving the Contents of the System Dump File After a System Failure ...	15-15
15.12	Freeing Dump Information from the Page File	15-16
15.13	Creating Page and Swap Files	15-17
15.13.1	Using AUTOGEN (Recommended Method)	15-17
15.13.2	Using SYSGEN	15-18
15.14	Installing Page and Swap Files	15-19
15.14.1	Installing Interactively	15-19
15.14.2	Installing in SYPAGSWPFILES.COM	15-20
15.15	Removing Page, Swap, and Dump Files	15-21
15.16	Changing Page, Swap, and Dump File Sizes	15-22
15.16.1	Using AUTOGEN (Recommended Method)	15-22
15.16.1.1	Controlling the Size of Page, Swap, and Dump Files in MODPARAMS.DAT	15-22
15.16.2	Using SWAPFILES.COM	15-25
15.16.3	Using SYSGEN	15-26
16	Performance Considerations	
16.1	Understanding Performance Management	16-1
16.2	Knowing Your Work Load	16-2
16.3	Choosing a Workload Management Strategy	16-3
16.4	Distributing the Work Load	16-4
16.5	Understanding System Tuning	16-4
16.6	Predicting When Tuning Is Required	16-5
16.7	Evaluating Tuning Success	16-6
16.8	Choosing Performance Options	16-6
16.9	Using INSTALL to Install Known Images	16-9
16.9.1	Understanding Images and Known Images	16-9
16.9.2	Understanding Known File Lists	16-10

16.9.3	Understanding Attributes You Can Assign to Known Images	16-10
16.9.4	Installing Images to Conserve Memory	16-11
16.9.5	Installing Images to Improve Image Performance	16-12
16.9.6	Installing Resident Images to Improve Performance (Alpha Only)....	16-12
16.9.7	Installing Images to Enhance Privileges of Images	16-13
16.9.7.1	Privileged Executable Images	16-14
16.9.7.2	Privileged Shareable Images	16-14
16.9.8	Installing Images to Allow Execution of Images Without Read Access	16-15
16.9.9	Determining Which Images to Install	16-15
16.9.10	Specifying File Names in INSTALL	16-16
16.9.11	Installing Images with INSTALL	16-16
16.9.12	Displaying Known Images with INSTALL	16-16
16.9.13	Defining Logical Names for Shareable Image Files	16-17
16.9.14	Removing Known Images	16-18

17 Testing the System with UETP

17.1	Overview	17-1
17.1.1	Understanding UETP	17-1
17.1.2	Summary of How to Use UETP	17-2
17.2	Preparing to Use UETP	17-3
17.2.1	Logging In	17-4
17.2.2	Using the SYSTEST Directories	17-4
17.3	Setting Up the Devices to Be Tested	17-4
17.3.1	Check Your Devices	17-5
17.3.2	System Disk Space Required	17-5
17.3.3	How UETP Works on Disks	17-5
17.3.4	Prepare Disk Drives	17-6
17.3.5	Magnetic Tape Drives	17-6
17.3.6	Tape Cartridge Drives	17-7
17.3.7	Compact Disc Drives	17-7
17.3.8	Optical Disk Drives	17-7
17.3.9	Terminals and Line Printers	17-8
17.3.10	Ethernet Adapters	17-8
17.3.11	DR11-W Data Interface (VAX Only)	17-8
17.3.12	DRV11-WA Data Interface (VAX Only)	17-9
17.3.13	DR750 or DR780 (DR32 Interface) (VAX Only)	17-9
17.3.14	Second LPA11-K Device	17-10
17.3.15	Devices That Are Not Tested	17-10
17.3.16	VMScluster Testing	17-10
17.3.17	Testing a Small-Disk System	17-11
17.3.18	DECnet for OpenVMS Phase	17-12
17.3.19	Vector Processors and the VVIEF (VAX Only)	17-12
17.4	Starting UETP	17-13
17.4.1	Running a Subset of Phases	17-13
17.4.2	Single Run Versus Multiple Passes	17-14
17.4.3	Defining User Load for Load Test	17-14
17.4.4	Report Formats	17-14
17.4.4.1	Long Report Format	17-15
17.4.4.2	Short Report Format	17-15
17.5	Stopping a UETP Operation	17-15
17.5.1	Using Ctrl/Y	17-15
17.5.2	Using DCL Commands	17-16

17.5.3	Using Ctrl/C	17-16
17.6	Troubleshooting: An Overview	17-16
17.6.1	Error Logging and Diagnostics	17-16
17.6.2	Interpreting UETP Output	17-17
17.6.3	Displaying Information on Your Screen	17-17
17.6.4	Example Screen Display (VAX Only)	17-18
17.6.5	Example Screen Display (Alpha Only)	17-19
17.6.6	Defining a Remote Node for UETP Ethernet Testing	17-19
17.6.7	Log Files	17-20
17.7	Troubleshooting: Possible UETP Errors	17-21
17.7.1	Summary of Common Failures	17-21
17.7.2	Wrong Quotas, Privileges, or Account	17-21
17.7.3	UETINIT01 Failure	17-23
17.7.4	UETVECTOR Failure (VAX Only)	17-25
17.7.5	Device Allocated or in Use by Another Application	17-25
17.7.6	Insufficient Disk Space	17-25
17.7.7	Incorrect Setup of a VMScluster System	17-26
17.7.8	Problems During the Load Test	17-27
17.7.9	DECnet for OpenVMS Error	17-28
17.7.10	Errors Logged but Not Displayed	17-29
17.7.11	No PCB or Swap Slots	17-29
17.7.12	No Keyboard Response or System Disk Activity	17-29
17.7.13	Lack of Default Access for the FAL Object	17-30
17.7.14	Bugchecks and Machine Checks	17-30
17.8	UETP Tests and Phases	17-30
17.8.1	Initialization Phase	17-31
17.8.2	Device Test Phase	17-31
17.8.2.1	How the Device Phase Works	17-31
17.8.2.2	Running a Single Device Test	17-32
17.8.2.3	Format of UETINIDEV.DAT	17-32
17.8.2.4	Running a Test in Loop Mode	17-33
17.8.2.5	Functions of Individual Device Tests	17-33
17.8.3	System Load Test Phase	17-34
17.8.4	DECnet for OpenVMS Test Phase	17-35
17.8.4.1	Environment	17-35
17.8.4.2	How the DECnet Phase Works	17-36
17.8.5	Cluster-Integration Test Phase	17-37

18 Getting Information About the System

18.1	Understanding System Log Files	18-1
18.2	Understanding Error Logging	18-2
18.3	Using the Error Formatter (ERRFMT)	18-3
18.3.1	Restarting the ERRFMT Process	18-3
18.3.2	Maintaining Error Log Files	18-3
18.3.3	Using ERRFMT to Send Mail	18-4
18.3.3.1	Enabling and Disabling ERRFMT to Send Mail	18-4
18.3.3.2	Sending Mail to Another User	18-5
18.4	Using the Error Log Utility (ERROR LOG)	18-5
18.4.1	Understanding the Error Log Utility (ERROR LOG)	18-5
18.4.2	Producing Error Log Reports	18-6
18.4.3	Producing a Full Error Log Report	18-6
18.4.4	Using Other Error Log Report Options	18-7

18.5	Using the DECEvent Event Management Utility (DECEvent) (Alpha Only)	18-8
18.5.1	Understanding DECEvent (Alpha Only)	18-8
18.5.2	Invoking and Exiting DECEvent (Alpha Only)	18-9
18.5.3	Using DECEvent Qualifiers (Alpha Only)	18-10
18.5.4	Using Additional DECEvent Commands (Alpha Only)	18-11
18.5.5	Producing DECEvent Reports (Alpha Only)	18-11
18.5.5.1	Producing a Full Report (Alpha Only)	18-11
18.5.5.2	Producing a Brief Report (Alpha Only)	18-12
18.5.5.3	Producing a Terse Report (Alpha Only)	18-13
18.5.5.4	Producing a Summary Report (Alpha Only)	18-14
18.5.5.5	Producing a Fast Error (FSTERR) Report (Alpha Only)	18-15
18.5.6	DECEvent Restrictions	18-15
18.6	Setting Up, Maintaining, and Printing the Operator Log File	18-16
18.6.1	Understanding the Operator Log File	18-16
18.6.2	Understanding OPCOM Messages	18-17
18.6.2.1	Initialization Messages	18-17
18.6.2.2	Device Status Messages	18-17
18.6.2.3	Terminal Enable and Disable Messages	18-17
18.6.2.4	User Request and Operator Reply Messages	18-19
18.6.2.5	Volume Mount and Dismount Messages	18-20
18.6.2.6	System Parameter Messages	18-20
18.6.2.7	Security Alarm Messages	18-21
18.6.2.8	Contents of an Operator Log File	18-21
18.6.3	Setting Up the Operator Log File	18-22
18.6.3.1	Creating a New Version of the Operator Log File	18-23
18.6.3.2	Specifying Logical Names	18-23
18.6.4	Maintaining the Operator Log File	18-24
18.6.5	Printing the Operator Log File	18-24
18.7	Using Security Auditing	18-25
18.7.1	Understanding Security Auditing	18-25
18.7.1.1	Security Audit Log File	18-26
18.7.1.2	Audit Log Files in Mixed-Version Clusters	18-26
18.7.2	Displaying Security Auditing Information	18-27
18.7.3	Delaying Startup of Auditing	18-27
18.7.4	Enabling Security Auditing for Additional Classes	18-28
18.7.5	Disabling Security Auditing	18-29
18.7.6	Enabling a Terminal to Receive Alarm Messages	18-29
18.7.7	Generating Security Reports	18-29
18.7.8	Creating a New Version of the Security Audit Log File	18-30
18.7.8.1	Creating a New Clusterwide Version of the Log File	18-30
18.7.8.2	Creating a New Node-Specific Version of the Log File	18-30
18.8	Monitoring Operating System Performance	18-30
18.8.1	Understanding the Monitor Utility (MONITOR)	18-31
18.8.1.1	MONITOR Classes	18-31
18.8.1.2	Display Data	18-32
18.8.1.3	Output Types	18-33
18.8.2	Invoking the Monitor Utility	18-33
18.8.3	Using Live Display Monitoring	18-34
18.8.4	Using Live Recording Monitoring	18-35
18.8.5	Using Concurrent Display and Recording Monitoring	18-36
18.8.6	Using Playback Monitoring	18-37
18.8.7	Using Remote Playback Monitoring	18-38
18.8.8	Rerecording Monitoring	18-39

18.8.9	Running MONITOR Continuously	18-39
18.8.9.1	Using the MONITOR.COM Procedure	18-40
18.8.9.2	Using the SUBMON.COM Procedure	18-41
18.8.9.3	Using the MONSUM.COM Procedure	18-42
18.8.10	Remote Monitoring in a Mixed-Version VMSccluster System	18-43

19 Tracking Resource Use

19.1	Understanding Accounting Files	19-1
19.2	Determining Which Resources Are Being Tracked	19-2
19.3	Controlling Which Resources Are Tracked	19-3
19.4	Starting Up a New Accounting File	19-3
19.5	Moving the Accounting File	19-3
19.6	Producing Reports of Resource Use	19-4
19.7	Setting Up Accounting Groups	19-5
19.8	Monitoring Disk Space	19-6

20 VMSccluster Considerations

20.1	Understanding VMSccluster Systems	20-1
20.1.1	Setting Up a VMSccluster Environment	20-2
20.1.2	Clusterwide System Management	20-3
20.2	Using DECams to Analyze Data	20-3
20.3	Using the Show Cluster Utility (SHOW CLUSTER)	20-4
20.3.1	Understanding the Show Cluster Utility	20-4
20.3.2	Beginning to Use SHOW CLUSTER Commands	20-6
20.3.2.1	Viewing Information That Is Off the Screen	20-6
20.3.2.2	Exiting from a Continuous Display	20-7
20.3.2.3	Using SHOW CLUSTER Qualifiers	20-7
20.3.3	Adding Information to a Report	20-8
20.3.4	Controlling the Display of Data	20-10
20.3.4.1	Entering Commands to Display Data	20-10
20.3.4.2	Removing Broadcast Messages	20-10
20.3.4.3	Refreshing the Screen	20-10
20.3.5	Formatting the Display of Data	20-11
20.3.5.1	Removing Information from a Report	20-11
20.3.5.2	Modifying Field and Screen Size	20-11
20.3.5.3	Moving a Report	20-12
20.3.5.4	Scrolling a Report	20-13
20.3.6	Creating a Startup Initialization File	20-14
20.3.7	Using Command Procedures Containing SHOW CLUSTER Commands	20-15
20.4	Understanding SYSMAN and VMSccluster Management	20-16
20.5	Using SYSMAN to Manage Security and System Time	20-16
20.5.1	Modifying the Group Number and Password	20-16
20.5.2	Modifying the System Time	20-17
20.5.2.1	Resetting System Time After January 1	20-18
20.6	Using the SYSMAN DO Command to Manage a VMSccluster	20-20

21 Network Considerations

21.1	Assigning Node Names	21-2
21.1.1	Syntax for Full Names	21-2
21.1.2	Considerations for Assigning Full Names	21-3
21.1.3	Setting Up a Node Name Strategy	21-3
21.2	Understanding DECnet for OpenVMS Networks	21-4
21.2.1	How an OpenVMS System Can Be Part of a Network	21-7
21.2.2	How Nodes Are Connected to the Network	21-8
21.2.3	Connecting Multiple Nodes to a SCSI Bus	21-9
21.2.4	Understanding the Configuration Database	21-10
21.2.5	How Your System Becomes a Node in the Network	21-11
21.3	Preparations for Joining a Network	21-12
21.4	Providing Security for Your Node	21-12
21.5	OpenVMS Support for TCP/IP Networking	21-13
21.5.1	Remote Terminal Service	21-13
21.5.2	Remote File Access	21-14
21.5.3	Remote Directory Listings	21-14
21.6	Managing a Network Node	21-14
21.6.1	Providing Host Services	21-15
21.6.2	Monitoring the Network	21-15
21.6.2.1	Using NCP Display Commands	21-15
21.6.2.2	NCP Counters	21-16
21.6.2.3	Using DECnet Event Logging	21-17
21.6.2.4	Using Other Software Tools	21-17
21.6.3	Testing the Network	21-17
21.6.4	Shutting Down and Restarting the Network	21-17

22 Managing the Local Area Network (LAN) Software

22.1	Understanding Local Area Networks	22-1
22.1.1	LAN Characteristics	22-2
22.1.1.1	Ethernet LANs	22-2
22.1.1.2	FDDI LANs	22-3
22.1.1.3	Token Ring LANs	22-3
22.1.2	LAN Addresses	22-3
22.2	Managing Local Area Networks	22-4
22.3	Understanding the LANACP LAN Server Process	22-5
22.3.1	Running the LANACP LAN Server Process	22-6
22.3.2	Stopping the LANACP LAN Server Process	22-6
22.4	Understanding the LANCP Utility	22-6
22.4.1	Invoking and Exiting LANCP	22-7
22.4.2	LANCP Commands	22-8
22.4.3	LANCP Miscellaneous Functions	22-9
22.5	Managing LAN Devices	22-9
22.5.1	Displaying System Devices	22-9
22.5.2	Displaying Device Parameters	22-10
22.5.3	Setting Device Parameters	22-12
22.5.4	Updating Device Firmware	22-14
22.6	Managing the LAN Device Databases	22-14
22.6.1	Displaying Devices in the LAN Device Databases	22-15
22.6.2	Entering Devices into the LAN Device Databases	22-15
22.6.3	Deleting Devices from the LAN Device Databases	22-17
22.7	Managing the LAN Node Databases	22-17
22.7.1	Displaying Nodes in the LAN Node Databases	22-18

22.7.2	Entering Nodes into the LAN Node Databases	22-18
22.7.3	Deleting Nodes from the LAN Node Databases	22-20
22.8	Understanding LAN MOP	22-20
22.8.1	Coexistence with DECnet MOP	22-20
22.8.2	Migrating from DECnet MOP to LAN MOP	22-20
22.8.3	Using CLUSTER_CONFIG_LAN.COM and LAN MOP	22-21
22.8.4	Sample Satellite Load	22-22
22.8.5	Cross-Architecture Booting	22-23
22.9	Managing the LAN MOP Downline Load Service	22-23
22.9.1	Enabling MOP Downline Load Service	22-23
22.9.2	Disabling MOP Downline Load Service	22-23
22.9.3	Displaying the Status and Counters Data	22-23
22.9.4	Displaying the Status and Counters Data for Individual Nodes	22-24
22.9.5	Clearing the Counters Data	22-25
22.9.6	OPCOM Messages	22-25
22.9.7	Load Trace Facility	22-25
22.9.8	MOP Console Carrier	22-26
22.9.9	MOP Trigger Boot	22-27

23 Managing InfoServer Systems

23.1	Understanding InfoServer Functions	23-1
23.1.1	Automatic Service Policies for Multiple Servers	23-4
23.1.2	High-Availability Feature to Reduce Service Interruptions	23-4
23.1.3	Support for X Terminal Clients	23-4
23.2	Understanding LASTport Protocols	23-5
23.2.1	LASTport Transport Protocol	23-5
23.2.2	LASTport/Disk Protocol	23-5
23.2.3	LASTport/Tape Protocol	23-6
23.3	Establishing a Server Management Session	23-6
23.3.1	Server Management Commands	23-7
23.4	Understanding InfoServer Client for OpenVMS Functions	23-8
23.5	Understanding LASTCP Utility Functions	23-9
23.5.1	Invoking and Exiting the LASTCP Utility	23-9
23.5.2	LASTCP Command Summary	23-10
23.5.3	Starting InfoServer Client for OpenVMS Software Automatically	23-10
23.5.4	Startup Restrictions: PATHWORKS and RSM	23-12
23.5.5	Startup Restrictions: SYSMAN	23-12
23.5.6	User Account Requirements	23-12
23.5.7	System Parameter MAXBUF Requirement	23-12
23.6	Understanding LADCP Utility Functions	23-12
23.6.1	Invoking and Exiting the LADCP Utility	23-13
23.6.2	LADCP Command Summary	23-13
23.6.3	Making InfoServer Devices Available Automatically	23-14

24 Managing the LAT Software

24.1	Understanding the LAT Protocol	24-1
24.1.1	How the LAT Protocol Works	24-2
24.1.2	Advantages of the LAT Protocol	24-2
24.2	Understanding the LAT Network	24-3

24.2.1	Service Nodes	24-4
24.2.1.1	Types of Services	24-4
24.2.1.2	Service Announcements	24-4
24.2.1.3	Print Requests	24-4
24.2.2	Terminal Server Nodes	24-5
24.2.2.1	Locating Service Nodes	24-5
24.2.2.2	Setting Up Connections	24-5
24.2.2.3	Servicing Nodes	24-5
24.2.3	Nodes Allowing Outgoing Connections	24-5
24.2.4	Sample LAT Configuration	24-6
24.2.5	LAT Relationship to VMSclusters and DECnet	24-7
24.2.5.1	LAT and DECnet Running on the Same Controller	24-7
24.2.5.2	LAT and DECnet Running on Different Controllers	24-8
24.3	Understanding the LATCP Utility	24-8
24.3.1	Invoking and Exiting LATCP	24-9
24.3.2	LATCP Commands	24-9
24.4	Starting Up the LAT Protocol	24-10
24.5	Customizing LAT Characteristics	24-12
24.5.1	Creating Additional Services	24-13
24.5.2	Setting Up Ports	24-14
24.5.2.1	Setting Up Printers	24-14
24.5.2.2	Setting Up Special Application Services	24-15
24.5.2.3	Setting Up Limited Services	24-15
24.5.3	Queuing Incoming Requests	24-15
24.5.4	Enabling Outgoing LAT Connections	24-16
24.5.5	Sample Edited LAT\$SYSTARTUP.COM Procedure	24-17
24.6	Managing the LATACP Database Size	24-19

25 Managing DECdtm Services

25.1	Understanding Transaction Logs	25-3
25.2	Planning Transaction Logs	25-3
25.2.1	Deciding the Size of a Transaction Log	25-3
25.2.2	Deciding the Location of a Transaction Log	25-4
25.3	Creating Transaction Logs	25-4
25.4	Monitoring Transaction Performance	25-6
25.5	Checking Whether a Transaction Log Is Too Small	25-8
25.6	Changing the Size of a Transaction Log	25-9
25.7	Moving a Transaction Log	25-11
25.8	Dismounting a Disk	25-14
25.9	Adding a Node	25-16
25.10	Removing a Node	25-18
25.11	Disabling DECdtm Services	25-20
25.12	Enabling DECdtm Services	25-21
25.13	Using DECdtm Services in a DECnet/OSI Network	25-22
25.13.1	Understanding the Configuration of a Transaction Group	25-22
25.13.2	Determining SCSNODE Name Uniqueness	25-24

26 Managing Special Processing Environments

26.1	Understanding Multiprocessing	26-2
26.1.1	Primary and Secondary Processors	26-2
26.1.2	Available and Active Sets	26-2
26.1.3	Processor Capabilities	26-2
26.2	Managing Symmetric Multiprocessing (SMP) Environments	26-3
26.2.1	Creating a Multiprocessing Environment	26-3
26.2.2	Monitoring a Multiprocessing Environment	26-3
26.3	Understanding Vector Processing	26-4
26.3.1	VAX Support for Vector Processing (VAX Only)	26-4
26.3.2	VAX Vector Instruction Emulation Facility (VVIEF) (VAX Only)	26-5
26.4	Managing the Vector Processing Environment (VAX Only)	26-5
26.4.1	Loading the Vector Processing Support Code (VAX Only)	26-5
26.4.2	Configuring a Vector Processing System (VAX Only)	26-6
26.4.3	Managing Vector Processes (VAX Only)	26-6
26.4.3.1	Adjusting System Resources and Process Quotas (VAX Only)	26-7
26.4.3.2	Distributing Scalar and Vector Resources Among Processes (VAX Only)	26-7
26.4.4	Restricting Access to the Vector Processor by Using ACLs (VAX Only)	26-8
26.4.5	Obtaining Information About a Vector Processing System (VAX Only)	26-8
26.4.5.1	DCL Lexical Functions F\$GETJPI and F\$GETSYI (VAX Only)	26-9
26.4.5.2	SHOW CPU/FULL Command (VAX Only)	26-9
26.4.5.3	SHOW PROCESS and LOGOUT/FULL Commands (VAX Only)	26-9
26.4.6	Loading the VAX Vector Instruction Emulation Facility (VVIEF) (VAX Only)	26-10

A Files-11 Disk Structure

A.1	Disk Concepts	A-1
A.1.1	Logical Organization of a Disk	A-1
A.1.2	Physical Organization of a Disk	A-2
A.2	Files-11 Structure	A-3
A.2.1	File Identification (FID)	A-4
A.2.2	ODS Directory Hierarchies	A-4
A.3	Reserved Files	A-4
A.3.1	Index File, INDEXF.SYS	A-5
A.3.1.1	Boot Block	A-6
A.3.1.2	Home Block	A-6
A.3.1.3	File Headers	A-7
A.3.2	Storage Bit Map File, BITMAP.SYS	A-8
A.3.3	Bad Block File, BADBLK.SYS	A-8
A.3.4	Master File Directory	A-8
A.3.5	Core Image File, CORIMG.SYS	A-9
A.3.6	Volume Set List File, VOLSET.SYS	A-9
A.3.7	Continuation File, CONTIN.SYS	A-9
A.3.8	Backup Log File, BACKUP.SYS	A-9
A.3.9	Pending Bad Block Log File, BADLOG.SYS	A-9
A.3.10	Quota File, QUOTA.SYS	A-9
A.3.11	Volume Security Profile, SECURITY.SYS	A-9
A.4	Files-11 ODS Level 1 Versus Level 2 (VAX Only)	A-10

B Tables of Time Differential Factors (TDFs)

Glossary

Index

Examples

14-1	Sample AUTOGEN Feedback Report	14-12
14-2	Sample AUTOGEN Command Procedure	14-23
18-1	Full Report Format (Alpha Only)	18-12
18-2	Brief Report Format (Alpha Only)	18-13
18-3	Terse Report Format (Alpha Only)	18-14
18-4	Summary Report Format (Alpha Only)	18-15
18-5	Fast Error (FSTERR) Report Format (Alpha Only)	18-15
18-6	Sample Operator Log File (SYS\$MANAGER:OPERATOR.LOG)	18-21
18-7	MONITOR.COM Procedure	18-40
18-8	SUBMON.COM Procedure	18-41
18-9	MONSUM.COM Procedure	18-42
20-1	SHOW CLUSTER Default Display	20-5
20-2	SHOW CLUSTER Display with CLUSTER Report	20-9

Figures

14-1	Old and New Parameter Values	14-7
14-2	SYSMAN Temporary, Active, and Current Parameter Values	14-26
14-3	SYSGEN Temporary, Active, and Current Parameter Values	14-31
21-1	Network Nodes, Circuits, and Lines	21-5
21-2	Example of a Small Local Area Network Configuration	21-6
21-3	Large Local Area Network Configuration	21-7
21-4	Highly Available Servers for Shared SCSI Access	21-10
23-1	InfoServer System Serving Clients	23-3
24-1	Sample LAT Network Configuration	24-7
25-1	Managing DECdtm Services	25-2
25-2	Transaction Group	25-23
A-1	File Extents	A-2
A-2	Tracks and Cylinders	A-3

Tables

14-1	Controlling AUTOGEN	14-10
14-2	AUTOGEN Phases	14-16
14-3	SYSMAN PARAMETERS Commands	14-25
14-4	SYSGEN Commands Used with System Parameters	14-30
15-1	Comparison of Physical and Selective Dumps	15-3
15-2	Paging and Swapping Terminology	15-4
15-3	Comparison of Physical and Selective Dump Files	15-10

15-4	Symbols for Controlling the Total Size of Page, Swap, or Dump File Space	15-24
15-5	Symbols for Controlling the Size of Individual Page and Swap Files	15-24
16-1	Attributes of Known Images	16-10
17-1	Device Tests (VAX Only)	17-34
17-2	Device Tests (Alpha Only)	17-34
18-1	System Log Files	18-2
18-2	Parts of the Error Logging Facility	18-2
18-3	Types of Events Reported in the Error Log File	18-5
18-4	Error Log Report Options	18-8
18-5	DECEvent Qualifiers (Alpha Only)	18-10
18-6	Event Classes Audited by Default	18-26
18-7	Types of MONITOR Classes	18-31
18-8	MONITOR Classes	18-32
18-9	MONITOR Command Procedures	18-39
18-10	Remote Monitoring Compatibility in a VMScluster System	18-43
20-1	Fields in Default SHOW CLUSTER Report	20-5
20-2	SHOW CLUSTER Qualifiers	20-8
20-3	Classes of Information Available in SHOW CLUSTER Reports	20-8
20-4	Fields in Sample CLUSTER Report	20-9
20-5	SYSMAN CONFIGURATION Commands	20-16
21-1	Sample System Names	21-3
21-2	DECnet for OpenVMS Network Terminology	21-4
21-3	Comparison of Volatile and Permanent Databases	21-11
21-4	Ways to Monitor the Network	21-15
21-5	NCP SHOW and LIST Commands	21-16
21-6	Network Monitoring Tools	21-17
22-1	Characteristics of LAN Media	22-2
22-2	LAN System Management Enhancements	22-4
22-3	LANACP System Logical Names	22-6
22-4	Functions of the LANCP Utility	22-7
22-5	Invoking the LANCP Utility	22-7
22-6	LANCP Commands	22-8
22-7	LAN Devices	22-9
22-8	SHOW DEVICE Command Qualifiers	22-10
22-9	SET DEVICE (parameters) Command Qualifiers	22-13
22-10	UPDATE DEVICE Command Qualifiers	22-14
22-11	LIST DEVICE and SHOW DEVICE Command Qualifiers	22-15
22-12	DEFINE DEVICE and SET DEVICE Command Qualifiers	22-16
22-13	DEFINE NODE and SET NODE Command Qualifiers	22-18
22-14	LIST NODE and SHOW NODE Command Qualifiers	22-24
22-15	CONNECT NODE Command Qualifiers	22-26
22-16	TRIGGER NODE Command Qualifiers	22-27
23-1	Summary of InfoServer Commands	23-7
23-2	Summary of LASTCP Commands	23-10
23-3	Summary of LADCP Commands	23-13

24-1	LATCP Commands	24-9
26-1	Settings of VECTOR_PROC System Parameter (VAX Only)	26-5
A-1	Reserved Files	A-5
A-2	Contents of Files-11 Index File	A-5
A-3	Areas of Data in File Headers	A-7
A-4	Limitations on Files-11 Structure Level 1 Volumes	A-10
B-1	TDFs for Europe	B-1
B-2	TDFs for North America	B-1
B-3	TDFs for Central and South America	B-2
B-4	TDFs for Asia	B-2
B-5	TDFs for the South Pacific	B-3
B-6	TDFs for Antarctica	B-3

Preface

Intended Audience

The intended audience for this manual is OpenVMS system managers.

Document Structure

The *OpenVMS System Manager's Manual: Tuning, Monitoring, and Complex Systems* consists of the following chapters:

- Chapter 14, Managing System Parameters
- Chapter 15, Managing System Page, Swap, and Dump Files
- Chapter 16, Performance Considerations
- Chapter 17, Testing the System with UETP
- Chapter 18, Getting Information About the System
- Chapter 19, Tracking Resource Use
- Chapter 20, VMScluster Considerations
- Chapter 21, Network Considerations
- Chapter 23, Managing InfoServer Systems
- Chapter 24, Managing the LAT Software
- Chapter 25, Managing DECdtm Services
- Chapter 26, Managing Special Processing Environments
- Appendix A, Files—11 Disk Structure
- Appendix B, Tables of Time Differential Factors (TDFs)
- Glossary

For more information about the structure of the *OpenVMS System Manager's Manual*, see Section 1.1 and the documentation map inside the front cover of this manual.

Related Documents

The following books are helpful when you use them in conjunction with the *OpenVMS System Manager's Manual*:

- *OpenVMS System Management Utilities Reference Manual*
- *OpenVMS User's Manual*
- *OpenVMS Software Overview*
- The current version of the *Upgrade and Installation Manual* for your system

- *OpenVMS Guide to System Security*
- *Guide to OpenVMS Performance Management*
- *VMScluster Systems for OpenVMS*
- The manuals in the networking kit of the OpenVMS Standard Documentation Set:
 - *DECnet for OpenVMS Guide to Networking*
 - *DECnet for OpenVMS Networking Manual*
 - *DECnet for OpenVMS Network Management Utilities*

Conventions

The name of the OpenVMS AXP operating system has been changed to OpenVMS Alpha. Any references to OpenVMS AXP or AXP are synonymous with OpenVMS Alpha or Alpha.

The following conventions are used to identify information specific to OpenVMS Alpha or to OpenVMS VAX:

Alpha

The Alpha icon denotes the beginning of information specific to OpenVMS Alpha.

VAX

The VAX icon denotes the beginning of information specific to OpenVMS VAX.

◆

The diamond symbol denotes the end of a section of information specific to OpenVMS Alpha or to OpenVMS VAX.

In this manual, every use of DECwindows and DECwindows Motif refers to DECwindows Motif for OpenVMS software.

The following conventions are also used in this manual:

Ctrl/*x*

A sequence such as Ctrl/*x* indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.

PF1 *x* or
GOLD

A sequence such as PF1 *x* or GOLD *x* indicates that you must first press and release the key labeled PF1 or GOLD and then press and release another key or a pointing device button.

GOLD key sequences can also have a slash (/), dash (-), or underscore (_) as a delimiter in EVE commands.

Return

In examples, a key name enclosed in a box indicates that you press a key on the keyboard. (In text, a key name is not enclosed in a box.)

...

Horizontal ellipsis points in examples indicate one of the following possibilities:

- Additional optional arguments in a statement have been omitted.
- The preceding item or items can be repeated one or more times.
- Additional parameters, values, or other information can be entered.

.
.
.

Vertical ellipsis points indicate the omission of items from a code example or command format; the items are omitted because they are not important to the topic being discussed.

()

In command format descriptions, parentheses indicate that, if you choose more than one option, you must enclose the choices in parentheses.

[]

In command format descriptions, brackets indicate optional elements. You can choose one, none, or all of the options. (Brackets are not optional, however, in the syntax of a directory name in an OpenVMS file specification or in the syntax of a substring specification in an assignment statement.)

{ }

In command format descriptions, braces indicate a required choice of options; you must choose one of the options listed.

boldface text

Boldface text represents the introduction of a new term or the name of an argument, an attribute, or a reason (user action that triggers a callback).

Boldface text is also used to show user input in Bookreader versions of the manual.

italic text

Italic text indicates important information, complete titles of manuals, or variables. Variables include information that varies in system messages (Internal error *number*), in command lines (/PRODUCER=*name*), and in command parameters in text (where *device-name* contains up to five alphanumeric characters).

struct

Monospace type in text identifies the following C programming language elements: keywords, the names of independently compiled external functions and files, syntax summaries, and references to variables or identifiers introduced in an example.

-

A hyphen in code examples indicates that additional arguments to the request are provided on the line that follows.

numbers

All numbers in text are assumed to be decimal unless otherwise noted. Nondecimal radices—binary, octal, or hexadecimal—are explicitly indicated.

...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...

...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...

...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...

...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...
...the ... of ...

...

...

...

Managing System Parameters

When your system is installed or upgraded, values of system parameters are automatically set by the command procedure SYS\$UPDATE:AUTOGEN.COM (AUTOGEN), which is supplied by Digital. Digital recommends you use AUTOGEN regularly to adjust the values for system parameters to fit your hardware configuration and your system's work load.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Converting your customized parameter settings for use with AUTOGEN	Section 14.3
Modifying system parameter values with AUTOGEN (recommended method)	Section 14.5
Controlling AUTOGEN's parameter settings with MODPARAMS.DAT	Section 14.5.1
Automating AUTOGEN reports	Section 14.6
Managing system parameters with SYSMAN	Section 14.7
Managing system parameters with SYSGEN	Section 14.8
Managing system parameters with a conversational boot	Section 14.9

This chapter explains the following concepts:

Concept	Section
System parameters	Section 14.1
Default, current, and active values of system parameters	Section 14.1.1
Pages and pagelets	Section 14.1.2
The recommended method for changing system parameter values	Section 14.2
The AUTOGEN command procedure	Section 14.4
AUTOGEN feedback	Section 14.4.1
The AUTOGEN feedback report (AGEN\$PARAMS.REPORT)	Section 14.4.2
AUTOGEN phases	Section 14.4.3
The AUTOGEN parameter file (MODPARAMS.DAT)	Section 14.4.4

14.1 Understanding System Parameters

The system uses values for **system parameters** to control how the system functions. System parameters control a wide range of system functions, including but not limited to the following:

- Memory management

Managing System Parameters

14.1 Understanding System Parameters

- Scheduling
- Security attributes
- System caches
- Windowing system choice
- Terminal configuration
- VAXcluster or VMScluster attributes

The *OpenVMS System Management Utilities Reference Manual* lists and describes each system parameter.

Your distribution kit provides **default values** for system parameters to allow you to boot any supported configuration. When your system is installed or upgraded, a command procedure supplied by Digital (SYS\$UPDATE:AUTOGEN.COM) executes to evaluate your hardware configuration, estimate typical work loads, and adjust the values of system parameters as needed.

Each system parameter has associated minimum and maximum values that define the scope of allowable values.

Parameter Types

System parameters can be one or more of the following types:

Type	Description
Dynamic	The value of a dynamic system parameter can be modified while the system is active by changing the <i>active</i> value in memory. In contrast, if you change the value of a parameter that is not dynamic, you must change the <i>current</i> value stored in the parameter file, and you must reboot the system for the changed value to take effect. For information on active and current values, see Section 14.1.1.
General	The value of a general parameter affects the creation and initialization of data structures at boot time.
Major	Major parameters are most likely to require modification.
Special	Special parameters are intended for use only by Digital. Change these parameters only if recommended by Digital personnel or in the installation guide or release notes of a Digital-supplied layered product.

Parameter Categories by Function

System parameters can be divided into the following categories, according to their function:

Category	Function
ACP	Parameters associated with file system caches and Files-11 XQP (extended QIO procedure) or ancillary control processes (ACPs). ¹
Cluster	Parameters that affect VAXcluster or VMScluster operation.
Job	Parameters that control jobs.
LGI	Parameters that affect login security.

¹Many ACP parameters are applicable only when Files-11 On-Disk Structure Level 1 disks are mounted or when an ACP is specifically requested during a mount command. In versions of the operating system before VAX VMS Version 4.0, a separate process, the Ancillary Control Process (ACP), performed file operations such as file opens, closes, and window turns. VAX VMS Version 4.0 introduced the XQP (extended QIO procedure), which allows every process on the system to perform these operations. For compatibility reasons, the names of the parameters have not changed.

Managing System Parameters

14.1 Understanding System Parameters

Category	Function
Multiprocessing	Parameters associated with symmetric multiprocessing.
PQL	Parameters associated with process creation limits and quotas.
RMS	Parameters associated with OpenVMS Record Management Services (RMS).
SCS	Parameters that control system communication services (SCS) and port driver operation. The parameters that affect SCS operation have the prefix SCS.
SYS	Parameters that affect overall system operation.
TTY	Parameters associated with terminal behavior.
User-defined	The following parameters can be user-defined: USERD1 (dynamic) USERD2 (dynamic) USER3 USER4

14.1.1 Default, Current, and Active Values

A system has several different sets of values for system parameters. The following table describes these values:

Value	Description
Default values	Values provided with the system to allow you to boot any supported configuration.
Current values	Values stored in the default parameter file on disk and used to boot the system. On VAX systems, the default parameter file is VAXVMSSYS.PAR. On Alpha systems, the default parameter file is ALPHAVMSSYS.PAR.
Active values	Values that are stored in memory and are used while the system is running. You can change the active value on a running system only for system parameters categorized as dynamic system parameters.
Values stored in other parameter files	For special purposes, you can create a parameter file other than the default parameter file that is used to store current values.

When the system boots, it reads the current values into memory, creating active values. An active value remains equal to the current value until you change either the active value or the current value.

When you execute the AUTOGEN command procedure through the SETPARAMS phase, it changes *current* values.

The System Management utility (SYSMAN) and the System Generation utility (SYSGEN) allow you to show and modify both *current* and *active* values. You use the USE and WRITE commands to specify which values you want to show or modify.

For more information about managing parameters with SYSMAN, see Section 14.7. For more information about managing parameters with SYSGEN, see Section 14.8.

Managing System Parameters

14.1 Understanding System Parameters

14.1.2 Pages and Pagelets

VAX

On VAX systems, the operating system allocates and deallocates memory for processes in units called **pages**. A page on a VAX system is 512 bytes. Some system parameter values are allocated in units of pages. ♦

Alpha

On Alpha systems, some system parameter values are allocated in units of pages, while others are allocated in units of **pagelets**.

A page on an Alpha system can be 8 kilobytes (KB) (8192 bytes), 16 KB, 32 KB, or 64 KB. A pagelet is a 512-byte unit of memory. One Alpha pagelet is the same size as one VAX page. On an Alpha computer with a page size of 8KB, 16 Alpha pagelets equal one Alpha page.

When reviewing parameter values, especially those parameters related to memory management, be sure to note the units required for each parameter. Section 14.7.2 and Section 14.8.2 explain how to show parameter values and their units of allocation. ♦

14.2 Recommended Method for Changing Parameter Values

Many system parameters can affect other parameters and the performance of the system. For this reason, Digital recommends you use the Digital-supplied command procedure `SYS$UPDATE:AUTOGEN.COM` (AUTOGEN) to manage system parameters. For information on AUTOGEN, see Section 14.4.

The System Management utility (SYSMAN) and the System Generation utility (SYSGEN) also allow you to manage system parameters. Although these utilities are not generally recommended for *changing* parameter values, you might want to use one of these utilities for the following reasons:

- To display system parameters and their values on a VAX or Alpha system
- To display system parameters and their values for systems on a VMScluster
- To temporarily modify a single parameter that has little effect on other parameters

Caution

If you change a parameter value with SYSMAN or SYSGEN, the value you set will be overridden or reset to the default value when you run AUTOGEN. To ensure that parameter changes are retained when you run AUTOGEN, you must add the parameter value to the AUTOGEN parameter file `MODPARAMS.DAT`. For more information, see Section 14.5.1.

If you currently use SYSMAN or SYSGEN to change parameters, and you have not added your customized parameter settings to `MODPARAMS.DAT`, follow the instructions in Section 14.3 to convert your customized parameter settings to `MODPARAMS.DAT` before running AUTOGEN.

14.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

14.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

Digital recommends you use the AUTOGEN command procedure to tune your system. For more information on AUTOGEN, see Section 14.4. If you use the System Management utility (SYSMAN) or the System Generation utility (SYSGEN) to modify system parameter values, and you do not include these changes in the AUTOGEN parameter file MODPARAMS.DAT, these changes will be overridden the next time you run AUTOGEN.

If you used SYSMAN or SYSGEN to change parameter values in the past, use the following procedure to convert your parameter settings to work with AUTOGEN. This procedure explains how to add your customized parameter settings to MODPARAMS.DAT so they will be retained when you run AUTOGEN.

Before performing this task, you should understand AUTOGEN, feedback, and the AUTOGEN parameter file MODPARAMS.DAT, as explained in Section 14.4.

1. Save the parameter values that the system is now using as follows:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE ACTIVE
SYSMAN> PARAMETERS WRITE SYS$SYSTEM:nodename_PARAMS_CURRENT.PAR
```

2. Write a listing of the active parameter values to an ASCII file named *nodename_PARAMS.OLD* as follows:

```
SYSMAN> PARAMETERS SHOW/ALL/OUTPUT=nodename_PARAMS.OLD
SYSMAN> PARAMETERS SHOW/SPECIAL/OUTPUT=nodename_PARAMS_SPECIAL.OLD
SYSMAN> EXIT
$ APPEND nodename_PARAMS_SPECIAL.OLD nodename_PARAMS.OLD
```

You will use this file in step 6.

3. Edit AUTOGEN's parameter file SYS\$SYSTEM:MODPARAMS.DAT to define symbols to specify values for the following:

- Parameter values that are not calculated by AUTOGEN, such as SCSNODE and SCSSYSTEMID. See the AUTOGEN description in the *OpenVMS System Management Utilities Reference Manual* for a table of the parameters calculated by AUTOGEN.
- Any parameter values that must be adjusted to suit your system work load, for example, GBLPAGES and GBLSECTIONS.

To specify a value, define symbols using the format MIN_parameter, MAX_parameter, or ADD_parameter rather than specifying an explicit value. For example:

```
$ EDIT SYS$SYSTEM:MODPARAMS.DAT

SCSNODE = "MYNODE" ! Not calculated by AUTOGEN
SCSSYSTEMID = 10001 ! Not calculated by AUTOGEN
MIN_GBLPAGES = 10000 ! Needed for MCS, BLISS32, and ADA
MIN_GBLSECTIONS = 600 ! Needed for MCS, BLISS32, and ADA
```

To help you track the changes you make in MODPARAMS.DAT, add comments to each line, preceded by the comment character (!). For information on defining symbols in MODPARAMS.DAT, see Section 14.5.1.

Managing System Parameters

14.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

4. Run AUTOGEN, but do *not* reboot. Use one of the following commands, depending on your system:

- If the system has run a typical work load for more than 24 hours since last booting:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS SETPARAMS FEEDBACK
```

The SAVPARAMS phase collects feedback information about resource use on the running system; this information is used by AUTOGEN. This command creates a feedback report named SYS\$SYSTEM:AGEN\$PARAMS.REPORT, which tells you about peak resource use.

- If you want to use a previously collected feedback file:

```
$ @SYS$UPDATE:AUTOGEN GETDATA SETPARAMS FEEDBACK
```

If you start from the GETDATA phase, AUTOGEN does not collect current feedback.

- If this is a new system (that is, it has no feedback) or the system has had little activity since last boot (for example, over the weekend) so there is no valid feedback file:

```
$@SYS$UPDATE:AUTOGEN GETDATA SETPARAMS CHECK_FEEDBACK
```

Use CHECK_FEEDBACK to let AUTOGEN determine whether the feedback is valid.

5. Write a listing of the new parameter values to an ASCII file as follows:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE CURRENT
SYSMAN> PARAMETERS SHOW /ALL /OUTPUT=nodename_PARAMS.NEW;
SYSMAN> PARAMETERS SHOW /SPECIAL /OUTPUT=nodename_PARAMS_SPECIAL.NEW;
SYSMAN> EXIT
$ APPEND nodename_PARAMS_SPECIAL.NEW; nodename_PARAMS.NEW;
```

6. Compare the old and new parameter values as follows:

```
$ DIFFERENCES/PARALLEL/OUTPUT=nodename_PARAMS.DIFF/MATCH=5 -
_$ nodename_PARAMS.OLD nodename_PARAMS.NEW
```

7. Print the differences file you created in step 6 (named in the format *nodename_PARAMS.DIFF*). Print the file on a 132-column line printer to make the output easier to read.
8. Compare the numbers in the two columns following each of the parameter name columns. The left column shows the old value; the right column shows the new value. Figure 14-1 illustrates sample output.

14.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

Figure 14–1 Old and New Parameter Values

Parameter Names	Old Values					New Values				
	GBLPAGES	77500	10000	512	–	GBLPAGES	81800	10000	512	–
SYSMWCNT	2400	500	40	1638		SYSMWCNT	2800	500	40	1638
INTSTKPAGES	4	4	1	–		INTSTKPAGES	4	4	1	–
BALSETCNT	250	16	4	819		BALSETCNT	250	16	4	819
WSMAX	32800	1024	60	20000		WSMAX	65500	1024	60	20000
NPAGEDYN	1944576	360000	16384	–		NPAGEDYN	3000000	360000	16384	–
NPAGEVIR	777328	1000000	16384	–		NPAGEVIR	12000000	1000000	16384	–
PAGEDYN	1516032	190000	10240	–		PAGEDYN	1780056	190000	10240	–
VIRTUALPAGECNT	150000	9216	512	100000		VIRTUALPAGECNT	270144	9216	512	100000
⋮	⋮					⋮	⋮			
⋮	⋮					⋮	⋮			

ZK-5175A-GE

- Make any adjustments in MODPARAMS.DAT using symbols prefixed by MIN_, MAX_, or ADD_. For example, if AUTOGEN calculated a smaller value for GBLPAGES, you might specify a minimum value for this parameter as follows:

```
MIN_GBLPAGES = 10000
```

If you originally specified a parameter value in MODPARAMS.DAT (in step 3) but the parameter has not been changed, verify the following:

- The parameter name is spelled correctly and completely (not abbreviated). In MODPARAMS.DAT, AUTOGEN sees parameter names as symbol assignments. AUTOGEN cannot equate a symbol to the corresponding system parameter unless it is spelled correctly. Look in AGEN\$FEEDBACK.REPORT for any error messages AUTOGEN might have written.
- The value is correct: count the digits and make sure no commas are present.
- The parameter occurs only once in MODPARAMS.DAT.
- The parameter is not commented out.

For most parameters, if the new value is greater than the old value, you can accept AUTOGEN's setting. If the new value is less than the old value, Digital recommends that you retain the old value because the system may not have been using that resource when running AUTOGEN.

For example, you might have used SYSMAN to increase GBLPAGES to 10,000 to accommodate layered products, but have not specified that change in MODPARAMS.DAT. AUTOGEN might calculate that the system needs only 5000 global pages. When you reboot after running AUTOGEN, not all of your layered products may be installed, and you might receive the system message GPTFULL, global page table full, indicating that the system needs more GBLPAGES.

- Repeat from Step 3 until you are satisfied with the new parameter values.

If necessary, make further changes in MODPARAMS.DAT, run AUTOGEN again, and verify the changes as before. Usually after this second pass of AUTOGEN, the parameter values will be stable and you can then reboot.

Managing System Parameters

14.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

11. Reboot. When you reboot, the system will use the new parameter values. Using AUTOGEN to reboot or rebooting right away are not necessary. However you must reboot before the system uses the new parameter values. If the system does not boot, perform a conversational boot and use the backup parameter file you created in step 1:

```
SYSBOOT> USE SYS$SYSTEM:nodename_PARAMS_CURRENT.PAR
SYSBOOT> CONTINUE
```

When you enter the CONTINUE command, the system boots with the parameter values you saved before running AUTOGEN.

After the system has booted, if you want to use the old parameter values you can enter the following commands:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE SYS$SYSTEM:nodename_PARAMS_CURRENT.PAR
SYSMAN> PARAMETERS WRITE CURRENT
SYSMAN> EXIT
```

12. Run AUTOGEN using feedback regularly to ensure that the resources of your system match your system work load. For information about running AUTOGEN using feedback, see Section 14.5.

14.4 Understanding the AUTOGEN Command Procedure

The AUTOGEN command procedure, SYS\$UPDATE:AUTOGEN.COM, is provided on your distribution kit, and runs automatically when your system is installed or upgraded to set appropriate values for system parameters. In addition, Digital recommends you run AUTOGEN when you want to reset values for system parameters or to resize page, swap, and dump files. The new values and file sizes take effect the next time the system boots.

AUTOGEN only calculates certain significant system parameters. See the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual* for a table of system parameters affected by AUTOGEN calculation.

When to Run AUTOGEN

Digital recommends running AUTOGEN in the following circumstances:

- During a new installation or upgrade. (This happens automatically as part of the installation or upgrade procedure.)
- Whenever your work load changes significantly.
- When you add an optional (layered) software product. See the specific product documentation for installation requirements. Certain layered products might require you to execute AUTOGEN to adjust parameter values and page and swap file sizes. (For information on using AUTOGEN to modify page and swap files, see Section 15.16.)
- When you install images with the /SHARED attribute; the GBLSECTIONS and GBLPAGES parameters might need to be increased to accommodate the additional global pages and global sections consumed.
- On a regular basis to monitor changes in your system's work load. You can automate AUTOGEN to regularly check feedback and recommend system parameter changes. Section 14.6 describes a batch-oriented command procedure that runs AUTOGEN in feedback mode on a regular basis and automatically sends the feedback report to an appropriate MAIL account.

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

- Periodically to provide adequate swapping file space. Use the FEEDBACK option and make sure the system has been up long enough (at least 24 hours) and that the load is typical. Also, make sure the SYS\$SYSTEM:MODPARAMS.DAT file does not contain a hardcoded SWAPFILE value, which prevents AUTOGEN from correctly sizing the swapping files.

AUTOGEN Operations

AUTOGEN executes in phases. Depending on which phases you direct it to execute, AUTOGEN performs some or all of the following operations:

- Collects the following types of data:
 - Feedback (from the running system)
 - The hardware configuration (from the system)
 - Parameter requirements supplied by you (from MODPARAMS.DAT)
 - Parameter requirements supplied by Digital
- Calculates appropriate new values for significant system parameters (listed in the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*).
- Creates a new installed image list.
- Calculates the sizes of system page, swap, and dump files.
- Adjusts the sizes of system page, swap, and dump files values of system parameter values, if necessary.
- Optionally shuts down and reboots the system.

Invoking AUTOGEN

To invoke AUTOGEN, enter a command in the following format at the DCL prompt:

```
@SYS$UPDATE:AUTOGEN [start-phase] [end-phase] [execution-mode]
```

Where:

<i>start-phase</i>	Is the phase where AUTOGEN is to begin executing. Section 14.4.3 lists the AUTOGEN phases.
<i>end-phase</i>	Is the phase where AUTOGEN is to complete executing. Section 14.4.3 lists the AUTOGEN phases.
<i>execution-mode</i>	Is one of the following: <ul style="list-style-type: none">• FEEDBACK—Use feedback.• NOFEEDBACK—Do not use feedback.• CHECK_FEEDBACK—Use feedback if it is valid. If feedback is invalid, ignore it, but continue executing through the end phase.• Blank (no execution mode specified)—Use feedback if it is valid. If it is not valid, quit before making any modifications.

For detailed information about invoking AUTOGEN, and the command line parameters you can specify, see the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*.

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

Controlling AUTOGEN Operations

Table 14–1 summarizes the methods for controlling AUTOGEN behavior.

Table 14–1 Controlling AUTOGEN

To Control...	Use This Method...
Which operations AUTOGEN is to perform	<p>Specify a start phase and an end phase when you invoke AUTOGEN.</p> <p>For detailed information about AUTOGEN phases, see the AUTOGEN section of the <i>OpenVMS System Management Utilities Reference Manual</i>.</p>
Parameter values set by AUTOGEN	<p>Specify values in the AUTOGEN parameter file MODPARAMS.DAT.</p> <p>You should periodically examine the results of calculations that AUTOGEN makes to determine whether AUTOGEN has drawn the correct conclusions about your hardware configuration and to be sure the system parameter values are appropriate for your workload requirements. If the values are not appropriate, adjust them by specifying desired values in MODPARAMS.DAT. For more information on MODPARAMS.DAT, see Section 14.4.4.</p>
AUTOGEN's use of feedback information	<p>Specify an execution mode when you invoke AUTOGEN.</p> <p>AUTOGEN can often improve system performance by using dynamic feedback gathered from the running system. However, feedback information is not always valid or appropriate. For more information, see Section 14.4.1.</p>

14.4.1 AUTOGEN Feedback

AUTOGEN feedback minimizes the need for you to modify parameter values or system file sizes. Instead, feedback allows AUTOGEN to automatically size the operating system based on your actual work load. **Sizing** is the process of matching the allocation of system resources (memory and disk space) with the workload requirements of your site.

Feedback is information, continuously collected by the operating system executive, about the amount of various resources the system uses to process its work load. The information is collected when exception events occur, so the collection does not affect system performance. When run in **feedback mode**, AUTOGEN analyzes this information and adjusts any related parameter values.

Note

When running AUTOGEN after making a major configuration change, specify **nofeedback** to assure the use of initial AUTOGEN settings. See Section 14.4).

AUTOGEN feedback affects the following resources (for a complete list of the affected system parameters, see the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*):

- Nonpaged pool
- Paged pool
- Lock resources

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

- Number of processes
- Global pages
- Global sections
- File system caches
- System logical name table sizes
- Page files
- Swap files

Feedback is gathered during AUTOGEN's SAVPARAMS phase and is written to the file SYS\$SYSTEM:AGEN\$FEEDBACK.DAT. This file is then read during the GETDATA phase. (See Section 14.4.3 for more information on AUTOGEN phases.)

Feedback is useful only if it accurately reflects the system's normal work load. For this reason, AUTOGEN performs some basic checks on the feedback and issues a warning message for either of the following conditions:

- The system has been up for less than 24 hours.
- The feedback is over 30 days old.

Whenever you modify the system (for example, a hardware upgrade, a change in the number of users, an optional product installation), you should operate in the new system environment for a period of time, and then execute AUTOGEN again starting from the SAVPARAMS phase.

VAX

On VAX systems, you can define the logical name AGEN\$FEEDBACK_REQ_TIME to specify, in hours, a minimum age required for feedback. For more information, see Section 14.5.2. ♦

When AUTOGEN runs, it displays whether feedback is used, as follows:

```
Feedback information was collected on 21-JAN-1995 14:00:08.53
Old values below are the parameter values at the time of collection.
The feedback data is based on 21 hours of up time.
Feedback information will be used in the subsequent calculations
```

See the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual* for a table of the system parameters affected by AUTOGEN feedback,

14.4.2 Feedback Report (AGEN\$PARAMS.REPORT)

Decides whether to use the system parameter values and system file sizes calculated by AUTOGEN. To help in your decision making, AUTOGEN generates a report file (SYS\$SYSTEM:AGEN\$PARAMS.REPORT) that includes the following information:

- All parameters and system files directly affected by the feedback
- Current values
- New values
- The feedback used in each parameter calculation
- Any user- or Digital-supplied modifications found in MODPARAMS.DAT
- Any advisory or warning messages displayed during AUTOGEN's operations

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

VAX

- On VAX systems, any user- or Digital-supplied modifications found in VMSPARAMS.DAT ♦

Alpha

- On Alpha systems, the parameters found during the GENPARAMS phase. ♦

Example 14-1 shows the contents of a sample AUTOGEN feedback report for a VAX system. On Alpha systems, the feedback report is similar but not identical to this example.

Suppression of Informational Messages

To suppress the display of informational messages, define the AGEN\$REPORT_NO_INFORMATIONALS logical to TRUE. Messages are entered in SYS\$SYSTEM:AGEN\$PARAMS.REPORT regardless of the value of AGEN\$REPORT_NO_INFORMATIONALS.

Example 14-1 Sample AUTOGEN Feedback Report

AUTOGEN Parameter Calculation Report on node: NODE22

This information was generated at 23-APR-1995 01:45:47.87

AUTOGEN was run from GETDATA to TESTFILES using FEEDBACK

** No changes will be done by AUTOGEN **

The values given in this report are what AUTOGEN would have set the parameters to.

Processing Parameter Data files

** WARNING ** - The system was up for less than 24 hours when the feedback information was recorded. This could result in feedback information that does not accurately reflect your typical work load.

Including parameters from: SYS\$SYSTEM:MODPARAMS.DAT

The following was detected within MODPARAMS.DAT

Please review immediately.

** INFORMATIONAL ** - Multiple MIN values found for MIN_CHANNELCNT.
Using MODPARAMS value (550) which is superseding OpenVMS value (255)

** INFORMATIONAL ** - Multiple MIN values found for MIN_SWPOUTPGCNT.
Using MODPARAMS value (1000) which is superseding OpenVMS value (500)

** INFORMATIONAL ** - Multiple MIN values found for MIN_PQL_DWSEXTENT.
Using MODPARAMS value (11000) which is superseding OpenVMS value (1024)

** INFORMATIONAL ** - Multiple MIN values found for MIN_PQL_MWSEXTENT.
Using MODPARAMS value (11000) which is superseding OpenVMS value (1024)

Feedback information was collected on 22-APR-1995 14:00:07.70

Old values below are the parameter values at the time of collection.

The feedback data is based on 13 hours of up time.

Feedback information will be used in the subsequent calculations

Parameter information follows:

MAXPROCESSCNT parameter information:

Feedback information.

Old value was 100, New value is 80

Maximum Observed Processes: 52

Information on VMS executable image Processing:

Processing SYS\$MANAGER:VMS\$IMAGES_MASTER.DAT

(continued on next page)

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

Example 14-1 (Cont.) Sample AUTOGEN Feedback Report

GBLPAGFIL parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 1024. The new value is 6024.
GBLPAGFIL has been increased by 5000.
GBLPAGFIL is not allowed to be less than 6024.

GBLPAGES parameter information:

Feedback information.
Old value was 43300, New value is 50000
Peak used GBLPAGES: 36622
Global buffer requirements: 6024

GBLSECTIONS parameter information:

Feedback information.
Old value was 400, New value is 400
Peak used GBLSECTIONS: 294
Override Information - parameter calculation has been overridden.
The calculated value was 350. The new value is 400.
GBLSECTIONS is not allowed to be less than 400.

LOCKIDTBL parameter information:

Feedback information.
Old value was 2943, New value is 3071
Current number of locks: 1853
Peak number of locks: 3200

LOCKIDTBL_MAX parameter information:

Feedback information.
Old value was 65535, New value is 65535

RESHASHTBL parameter information:

Feedback information.
Old value was 1024, New value is 1024
Current number of resources: 957

MSCP_LOAD parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 1. The new value is 0.
MSCP_LOAD has been disabled by a hard-coded value of 0.

MSCP_BUFFER parameter information:

Feedback information.
Old value was 128, New value is 128
MSCP server I/O rate: 0 I/Os per 10 sec.
I/Os that waited for buffer space: 0
I/Os that fragmented into multiple transfers: 0

SCSCONNCNT parameter information:

Feedback information.
Old value was 5, New value is 5
Peak number of nodes: 1
Number of CDT allocation failures: 0

SCSRESCNT parameter information:

Feedback information.
Old value was 300, New value is 300
RDT stall count: 0

SCSBUFFCNT parameter information:

Feedback information.
Old value was 512, New value is 512
CIBDT stall count: 0

(continued on next page)

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

Example 14-1 (Cont.) Sample AUTOGEN Feedback Report

NPAGEDYN parameter information:

Feedback information.

Old value was 686592, New value is 783360

Maximum observed non-paged pool size: 815616 bytes.

Non-paged pool request rate: 47 requests per 10 sec.

LNMSHASHTBL parameter information:

Feedback information.

Old value was 1024, New value is 1024

Current number of shareable logical names: 1194

ACP_DIRCACHE parameter information:

Feedback information.

Old value was 88, New value is 88

Hit percentage: 99%

Attempt rate: 0 attempts per 10 sec.

ACP_DINDXCACHE parameter information:

Feedback information.

Old value was 25, New value is 25

Hit percentage: 97%

Attempt rate: 1 attempts per 10 sec.

ACP_HDRCACHE parameter information:

Feedback information.

Old value was 88, New value is 106

Hit percentage: 98%

Attempt rate: 17 attempts per 10 sec.

ACP_MAPCACHE parameter information:

Feedback information.

Old value was 8, New value is 8

Hit percentage: 2%

Attempt rate: 4 attempts per 10 sec.

PAGEDYN parameter information:

Feedback information.

Old value was 521728, New value is 542208

Current paged pool usage: 304160 bytes.

Paged pool request rate: 1 requests per 10 sec.

PFRATL parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 0. The new value is 1.

PFRATL has been disabled by a hard-coded value of 1.

WSDEC parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 35. The new value is 19.

WSDEC has been disabled by a hard-coded value of 19.

MPW_LOLIMIT parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 120. The new value is 2100.

MPW_LOLIMIT is not allowed to be less than 2100.

MPW_HILIMIT parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 1310. The new value is 4500.

MPW_HILIMIT is not allowed to be less than 4500.

(continued on next page)

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

Example 14-1 (Cont.) Sample AUTOGEN Feedback Report

LONGWAIT parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 30. The new value is 10.

LONGWAIT has been disabled by a hard-coded value of 10.

WSMAX parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 8200. The new value is 12000.

WSMAX is not allowed to be less than 12000.

PROCSECTCNT parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 32. The new value is 40.

PROCSECTCNT is not allowed to be less than 40.

PQL_DWSEXTENT parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 400. The new value is 11000.

PQL_DWSEXTENT is not allowed to be less than 11000.

PQL_MWSEXTENT parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 2048. The new value is 11000.

PQL_MWSEXTENT is not allowed to be less than 11000.

VAXCLUSTER parameter information:

Override Information - parameter calculation has been overridden.

The calculated value was 1. The new value is 0.

VAXCLUSTER has been disabled by a hard-coded value of 0.

Page, Swap, and Dump file calculations

Page and Swap file calculations.

PAGEFILE1_SIZE parameter information:

Feedback information.

Old value was 45200, New value is 50500

Maximum observed usage: 25265

PAGEFILE1_SIZE will be modified to hold 50500 blocks

PAGEFILE2_SIZE parameter information:

Feedback information.

Old value was 154000, New value is 194400

Maximum observed usage: 97175

PAGEFILE2_SIZE will be modified to hold 194400 blocks

** WARNING ** - The disk on which PAGEFILE2 resides would be over 95% full if it were modified to hold 194400 blocks.

NODE22\$DKA300:[SYSTEM_FILES]PAGEFILE.SYS will not be modified.

NODE22\$DKA300:[SYSTEM_FILES]PAGEFILE.SYS will remain at 154002 blocks.

SWAPFILE1_SIZE parameter information:

Feedback information.

Old value was 15000, New value is 15000

Maximum observed usage: 14280

Override Information - parameter calculation has been overridden.

The calculated value was 21400. The new value is 15000.

SWAPFILE1_SIZE is not allowed to exceed 15000.

SWAPFILE1 will not be modified.

(continued on next page)

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

Example 14-1 (Cont.) Sample AUTOGEN Feedback Report

```
SWAPFILE2_SIZE parameter information:
Feedback information.
  Old value was 50000, New value is 26300
  Maximum observed usage: 1680
SWAPFILE2_SIZE will be modified to hold 26300 blocks

** WARNING ** - The disk on which SWAPFILE2 resides would be
over 95% full if it were modified to hold 26300 blocks.
NODE22$DKA300:[SYSTEM_FILES]SWAPFILE.SYS will not be modified.
NODE22$DKA300:[SYSTEM_FILES]SWAPFILE.SYS will remain at 50001 blocks.

Dumpfile calculations:
No dump file modifications would have been made.
Dumpfile will remain at 34116 blocks.
```

14.4.3 AUTOGEN Phases

When you invoke AUTOGEN, you specify a start phase and an end phase for AUTOGEN to execute. AUTOGEN executes all phases from the start phase to the end phase. Depending on the start phase and end phase you specify, AUTOGEN can execute any of the following phases, in the order shown in Table 14-2.

Table 14-2 AUTOGEN Phases

Phase	Description
SAVPARAMS	Saves dynamic feedback from the running system.
GETDATA	Collects all data to be used in AUTOGEN calculations.
GENPARAMS	Generates new system parameters; creates the installed image list.
TESTFILES	Displays the system page, swap, and dump file sizes calculated by AUTOGEN (cannot be used as a start phase).
GENFILES	Generates new system page, swap, and dump files if appropriate (cannot be used as a start phase).
SETPARAMS	Runs SYSMAN to set the new system parameters in the default parameter file, saves the original parameters, and generates a new parameter file, AUTOGEN.PAR. On VAX systems, the default parameter file is VAXVMSSYS.PAR. The original parameters are saved in the file VAXVMSSYS.OLD. On Alpha systems, the default parameter file is ALPHAVMSSYS.PAR. The original parameters are saved in the file ALPHAVMSSYS.OLD.
SHUTDOWN	Prepares the system to await a manual reboot.
REBOOT	Automatically shuts down and reboots the system.
HELP	Displays help information to the screen.

For detailed information about each AUTOGEN phase and the files affected by each phase, see the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*.

Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

14.4.4 AUTOGEN Parameter File (MODPARAMS.DAT)

AUTOGEN reads a parameter file named MODPARAMS.DAT during the GETDATA phase. You can add commands to this file to control the system parameter values and file sizes that AUTOGEN sets. You can use MODPARAMS.DAT to do the following:

Operation	For More Information
Increase the value of any numeric system parameter	Section 14.5.1.1
Set a minimum value for a numeric system parameter	Section 14.5.1.2
Set a maximum value for a numeric system parameter	Section 14.5.1.3
Specify an absolute value for a system parameter	Section 14.5.1.4
Include an external parameter file	Section 14.5.3
Specify sizes of page, swap, and dump files	Section 15.16.1.1
Define the number of VMScluster nodes	Section 14.5.1.5
†Define the number of Ethernet adapters	Section 14.5.1.6
Preset parameter values before adding memory	Section 14.5.1.7
Specify an alternate default startup command procedure	Section 4.4.2
†VAX specific	

To help track changes you make to MODPARAMS.DAT, make sure you add comments, preceded by the comment character (!), each time you change the file.

Caution

The recommended method of changing system parameters and system file sizes is to edit MODPARAMS.DAT to include parameter settings. If you change a system parameter value or file size using SYSMAN, SYSGEN, or a conversational boot, and you do not specify the value in MODPARAMS.DAT, AUTOGEN will recalculate the value or file size the next time it runs. For more information, see Section 14.5.1.

Example

The following example shows the contents of a sample MODPARAMS.DAT file:

```
!
! ***** A Sample MODPARAMS.DAT for Node NODE22 *****
!
! MODPARAMS.DAT for "NODE22"
! REVISED: 04/29/95 -CHG- Upped GBLPAGES to account for ADA.
!
SCSNODE          = "NODE22"      ! This is not calculated by AUTOGEN.
SCSSYSTEMID      = 19577         ! This is not calculated by AUTOGEN.
TTY_DEFCHAR2     = %X0D34       ! This is not calculated by AUTOGEN.
ADD_ACP_DIRCACHE= 150           ! Hit rate was only 65% on directory cache.
MIN_PAGEDYN      = 500000       ! PAGEDYN must be at least 1/2 Mbyte to
                                ! account for a large number of logical names.
!
MAX_PAGEFILE1_SIZE = 15000      ! Maximum size for primary page.
MAX_SWAPFILE      = 5000       ! Maximum size for swap file space.
MAX_DUMPFILE      = 32768      ! Maximum size for dump file space.
```


Managing System Parameters

14.4 Understanding the AUTOGEN Command Procedure

```
ADD_GBLPAGES      = 425+507+157    ! Account for MCS, BLISS32 and ADA.
ADD_GBLSECTIONS   = 4 + 5 + 2      ! Account for MCS, BLISS32 and ADA.
VIRTUALPAGECNT    = 144264         ! So that we can read MONSTR's 68Mb dumps.
!
! end of MODPARAMS.DAT for NODE22
```

14.5 Modifying System Parameters with AUTOGEN

The recommended method of modifying system parameters is to execute AUTOGEN in two passes, as follows:

1. First pass—Execute AUTOGEN using the following command:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS GENPARAMS
```

This command instructs AUTOGEN to do the following:

- Save the current feedback
- Gather all of the information required for the calculations
- Calculate the system parameter values
- Generate the feedback report
- Write the information to SETPARAMS.DAT

Review the input to the calculations (PARAMS.DAT), the output from the calculations (SETPARAMS.DAT), and the report generated (AGEN\$PARAMS.REPORT).

If you are not satisfied with the parameter settings, modify parameter values by editing MODPARAMS.DAT as explained in Section 14.5.1. Then reexecute AUTOGEN from the GETDATA phase.

When you are satisfied with the contents of SETPARAMS.DAT, go on to step 2.

2. Second pass—Execute AUTOGEN a second time using the following command:

```
$ @SYS$UPDATE:AUTOGEN GENPARAMS REBOOT
```

This AUTOGEN command runs SYSMAN to update the new system parameter values and installs them on the system when it is rebooted.

14.5.1 Controlling AUTOGEN's Parameter Settings with MODPARAMS.DAT

If, after examining the AGEN\$PARAMS.REPORT or SETPARAMS.DAT file, you decide to correct hardware configuration data or modify system parameter values chosen by AUTOGEN, edit the MODPARAMS.DAT file as described in this section to manually specify parameter values.

Caution

Always edit MODPARAMS.DAT to specify values for parameters. Do not edit PARAMS.DAT; modifying the contents of this file might prevent AUTOGEN from operating correctly.

For information on editing MODPARAMS.DAT to control sizes of page, swap, and dump files, see Section 15.16.1.1.

Managing System Parameters

14.5 Modifying System Parameters with AUTOGEN

You can define symbols in MODPARAMS.DAT using the following formats to control parameter values:

Control Method	Symbol Format	For More Information
Increase a value by a specified amount	ADD_*	Section 14.5.1.1
Specify a minimum value	MIN_*	Section 14.5.1.2
Specify a maximum value	MAX_*	Section 14.5.1.3
Specify an absolute value	Parameter name	Section 14.5.1.4

When defining symbols in MODPARAMS.DAT, make sure of the following:

- The value is correct and valid for the parameter. Count the digits. Do not use commas.
- The symbol occurs only once in MODPARAMS.DAT.
- The symbol value is not commented out.
- The symbol name is spelled correctly and completely (not abbreviated).

Caution

When AUTOGEN reads MODPARAMS.DAT or any other parameter file, it checks to determine if the symbol names specified in the file are valid. If they are not, AUTOGEN writes a warning message to AGEN\$PARAMS.REPORT. However, AUTOGEN checks only the symbol name; it does not check the validity of the value specified for the symbol.

If a value is invalid, the line is *not* ignored. AUTOGEN attempts to use the specified value.

A symbol is not checked if it is specified in a line that contains a DCL expression other than the symbol assignment (=). For example, AUTOGEN does not check the validity of a symbol name specified in a line with the DCL IF statement. Instead, AUTOGEN writes a warning message to AGEN\$PARAMS.REPORT.

To help track changes you make to MODPARAMS.DAT, make sure you add comments preceded by the comment character (!) each time you change the file.

14.5.1.1 Increasing a Value with the ADD_ Prefix

Use the ADD_ prefix to increase the value of any NUMERIC parameter. The new values are updated in subsequent AUTOGEN calculations during the GENPARAMS phase. The following example demonstrates the use of the ADD_ prefix:

```
ADD_GBLPAGES=500
ADD_NPAGEDYN=10000
```

An ADD_ parameter record for a parameter that AUTOGEN calculates will add the value to AUTOGEN's calculations. An ADD_ parameter record for a parameter that AUTOGEN does not calculate will add the value to the parameter's default (not current) value. (See the AUTOGEN section on the *OpenVMS System Management Utilities Reference Manual* for a table of parameters affected by AUTOGEN.)

Managing System Parameters

14.5 Modifying System Parameters with AUTOGEN

Note

The ADD_ value is added to the calculated value once, and does not accumulate with successive runs for feedback calculations.

Typically, you would not use the ADD_ prefix for modifying parameters that are calculated by the feedback mechanism, because the feedback results should accurately reflect your work load. However, if you do use the ADD_ prefix with feedback, be aware that AUTOGEN will add a value only once if AUTOGEN is run to the SETPARAMS phase or beyond. If you wish to maintain a minimum level above AUTOGEN's calculation, use the MIN_ prefix.

14.5.1.2 Specifying a Minimum Value with the MIN_ Prefix

Use the MIN_ prefix if you do not want AUTOGEN to set a parameter below a specified value. MIN_ refers to the minimum value to which a parameter can be set by AUTOGEN.

```
MIN_PAGEDYN = 400000
```

14.5.1.3 Specifying a Maximum Value with the MAX_ Prefix

Use the MAX_ prefix if you do not want AUTOGEN to set a parameter above a specified value. MAX_ refers to the maximum value to which a parameter can be set by AUTOGEN.

```
MAX_PAGEDYN = 400000
```

14.5.1.4 Specifying an Absolute Value

Use this method to specify a value for a parameter that AUTOGEN does not calculate. (See the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual* for a table of the system parameters modified in AUTOGEN calculations.)

Note

Digital strongly recommends that you use this method only for parameters that describe the system environment (for example, SCSNODE and SCSYSTEMID). For the parameters that AUTOGEN calculates, specifying a value with this method disables AUTOGEN's calculations. Instead of specifying an absolute value, use one of the following methods:

- Specify a minimum value with the MIN_ prefix (see Section 14.5.1.2)
 - Specify a maximum value with the MAX_ prefix (see Section 14.5.1.3)
 - Increase the value with the ADD_ prefix (see Section 14.5.1.1)
-

To specify an absolute parameter value, add an assignment statement in the following format to MODPARAMS.DAT:

```
parameter = parameter-value ! comment
```

For example, the following command assigns the node name BIGVAX to the SCSNODE parameter:

```
SCSNODE = "BIGVAX" ! the node name
```


Managing System Parameters

14.5 Modifying System Parameters with AUTOGEN

14.5.1.5 Defining the Number of VAXcluster Nodes (VAX Only)

VAX

In a VAXcluster environment, use the `NUM_NODES` symbol to prevent temporary changes in VAXcluster membership from affecting AUTOGEN's calculation of VAXcluster-related parameter values. Define the `NUM_NODES` symbol in `MODPARAMS.DAT` to specify the number of nodes that are to run in the VAXcluster. AUTOGEN uses this value to set parameters that are affected by the number of VAXcluster nodes.

For example, you might include the following line in `MODPARAMS.DAT`:

```
NUM_NODES = 30 ♦
```

14.5.1.6 Defining the Number of Ethernet Adapters (VAX Only)

VAX

In a VAXcluster environment, define the `NUM_ETHERADAPT` symbol in `MODPARAMS.DAT` to specify the total number of Ethernet adapters in the VAXcluster system. For example, you might include the following line in `MODPARAMS.DAT`:

```
NUM_ETHERADAPT = 40 ♦
```

14.5.1.7 Presetting Parameter Values Before Adding Memory (VAX Only)

VAX

On VAX systems, if you are planning to upgrade your system hardware by adding a large amount (512 MB or more) of memory, you might want to preset your system parameters to values appropriate for the additional memory. Presetting your system parameters minimizes the possibility of memory upgrade problems caused by inappropriate parameter values.

How to Perform This Task

Perform the following steps:

1. Add a line to `SYS$SYSTEM:MODPARAMS.DAT` in the following format:

```
MEMSIZE = total-number-of-pages-of-memory-after-upgrade
```

For example:

```
MEMSIZE = 2048 * 1024 ! (2048 page per MB * 1GB of memory)
```

2. Run AUTOGEN to the `SETPARAMS` phase.
3. Perform the hardware upgrade to add the additional memory.
4. Edit `MODPARAMS.DAT` to remove the line added in step 1. ♦

14.5.1.8 Overriding Parameters Related to DECnet

To override AUTOGEN's observations regarding the presence (or absence) of DECnet, set the `MODPARAMS.DAT` parameter `LOAD_DECNET_IMAGES` to `TRUE` (or `FALSE`). Controlling the setting is useful for sites that have no synchronous network hardware but want to run asynchronous DECnet.

14.5.2 Specifying a Minimum Required Age for Feedback (VAX Only)

VAX

On VAX systems, AUTOGEN feedback is useful only when a system has been running long enough to accurately reflect the system's normal work load. By default, AUTOGEN uses feedback if the data is older than 24 hours. On VAX systems, you can define the logical name `AGEN$FEEDBACK_REQ_TIME` to specify, in hours, a different minimum age required for feedback. AUTOGEN uses this value to determine whether the feedback is to be used.

Managing System Parameters

14.5 Modifying System Parameters with AUTOGEN

For example, you might define the logical name as follows, to indicate that AUTOGEN should use feedback if it is older than 19 hours:

```
$ DEFINE/SYSTEM AGEN$FEEDBACK_REQ_TIME 19
```

To define this logical name each time the system starts up, add this command to SYLOGICALS.COM. ♦

14.5.3 Including an External Parameter File in MODPARAMS.DAT

You can include external parameter files in MODPARAMS.DAT. For example, you might want to set a system parameter to the same value on all nodes in a VAXcluster or VMScluster; you might also want to specify node-specific values for other system parameters. You could specify the cluster-common values in a separate cluster-common file and include this cluster-common file in the MODPARAMS.DAT file on each system in the VAXcluster or VMScluster.

To include a parameter file, place a command in the following format in MODPARAMS.DAT, or in any parameter file that is included in MODPARAMS.DAT:

```
AGEN$INCLUDE_PARAMS full-directory-spec:filename
```

Example

To include a cluster-common parameter file named CLUSTERPARAMS.DAT, create a common parameter file with the following name:

```
SYS$COMMON:[SYSEXE]CLUSTERPARAMS.DAT
```

Add the following line in the MODPARAMS.DAT file in the system-specific directory of each VMScluster system:

```
AGEN$INCLUDE_PARAMS SYS$COMMON:[SYSEXE]CLUSTERPARAMS.DAT
```

14.6 Automating AUTOGEN Reports

Digital recommends you create a batch-oriented command procedure to automatically run AUTOGEN on a regular basis and send the resulting feedback reports to an appropriate MAIL account. Example 14–2 provides a sample command procedure.

Note

This command procedure runs AUTOGEN only to recommend system parameter values and send you a report. It does not run AUTOGEN to change system parameters or reboot the system. If, after reviewing the report, you decide to change system parameters, follow the instructions in Section 14.6.1.

The command procedure in Example 14–2 runs two passes of AUTOGEN. On the first pass, AUTOGEN runs during peak workload times to collect data on realistic system work loads. This pass does not degrade system performance. On the second pass, AUTOGEN runs during off-peak hours to interpret the data collected in the first stage.

The procedure sends the resulting report, contained in the file AGEN\$PARAMS.REPORT, to the SYSTEM account. Review this report on a regular basis to see whether the load on the system has changed.

Managing System Parameters

14.6 Automating AUTOGEN Reports

Example 14-2 shows a sample command procedure. Use this procedure only as an example; create a similar command procedure as necessary to meet the needs of your configuration.

Example 14-2 Sample AUTOGEN Command Procedure

```
$ BEGIN$: ! ++++++ AGEN_BATCH.COM ++++++
$ on warning then goto error$
$ on error then goto error$
$ on severe_error then goto error$
$ on control_y then goto error$
$!
$! Setup process
$!
$! Set process information
$ set process/priv=all/name="AUTOGEN Batch"
$! Keep log files to a reasonable amount
$ purge/keep=5 AGEN_Batch.log
$ time = f$time() ! Fetch current time
$ hour = f$integer(f$cvtime(time,,"hour")) ! Get hour
$ today = f$cvtime(time,,"WEEKDAY") ! Get Day of the week
$ if f$integer(f$cvtime(time,,"minute")) .ge. 30 then hour = hour + 1
$!
$! Start of working day...
$!
$ 1AM$:
$ if hour .le. 2
$ then
$ next_time = "today+0-14"
$ gosub submit$ ! Resubmit yourself
$ set noon
$!
$! Run AUTOGEN to TESTFILES using the parameter values collected earlier
$! in the day (i.e., yesterday at 2:00pm)
$ if today .eqs. "Tuesday" .OR. today .eqs. "Thursday" .OR. -
today .eqs. "Saturday"
$ then
$ @sys$update:autogen GETDATA TESTFILES feedback ②
$ mail/sub="AUTOGEN Feedback Report for system-name" -
sys$system:agen$params.report system ③
$ ! Clean up
$ purge/keep=7 sys$system:agen$feedback.report ④
$ purge/keep=7 sys$system:agen$feedback.dat
$ purge/keep=7 sys$system:params.dat
$ purge/keep=7 sys$system:autogen.par
$ purge/keep=7 sys$system:setparams.dat
$ purge/keep=7 sys$system:agen$addhistory.tmp
$ purge/keep=7 sys$system:agen$addhistory.dat
$ endif
$ goto end$
$ endif
```

(continued on next page)

Managing System Parameters

14.6 Automating AUTOGEN Reports

Example 14-2 (Cont.) Sample AUTOGEN Command Procedure

```
$!  
$ 2PM$:  
$   if hour .le. 15  
$       then  
$           next_time = "today+0-17"  
$           gosub submit$  
$           if today .eqs. "Monday" .OR. today .eqs. "Wednesday" .OR. -  
$               today .eqs. "Friday"  
$               then  
$                   @sys$update:autozen SAVPARAMS SAVPARAMS feedback ❶  
$                   endif  
$               goto end$  
$           endif  
$!  
$ 5PM$:  
$   if hour .le. 18  
$       then  
$           next_time = "tomorrow+0-1"  
$           gosub submit$  
$           endif  
$!  
$! End of working day...  
$!  
$ END$:      ! ----- BATCH.COM -----  
$ exit  
$!++  
$! Subroutines  
$!--  
$!  
$ SUBMIT$:  
$   submit/name="AGEN_Batch"/restart/noprint - ❺  
$   /log=AGEN_batch.log -  
$   /queue=sys$batch/after="'next_time'" sys$system:AGEN_batch.com  
$   return  
$!++  
$! Error handler  
$!--  
$ ERROR$:  
$   mail/sub="AGEN_BATCH.COM - Procedure failed." _nl: system  
$   goto end$
```

The commands in this procedure perform the following tasks:

- ❶ Executes the first pass of AUTOGEN during peak workload times to collect data on realistic work loads. This command runs a very fast image so it does not degrade system response.
- ❷ Executes the second pass of AUTOGEN during off-peak hours to interpret the data collected in the first pass.
- ❸ Mails the resulting report file named AGEN\$PARAMS.REPORT to the SYSTEM account.
- ❹ Cleans up the files created.
- ❺ Resubmits the command procedure.

14.6.1 Changing Parameter Values After Reviewing AUTOGEN Reports

If the command procedure report described in the previous section shows AUTOGEN's calculations are different from the current values, correct the tuning by executing AUTOGEN with one of the two following commands:

- If the system can be shut down and rebooted immediately, execute the following command:

```
$ @SYS$UPDATE:AUTOGEN GETDATA REBOOT FEEDBACK
```

- If the system cannot be shut down and rebooted immediately, execute the following command to reset the system parameters:

```
$ @SYS$UPDATE:AUTOGEN GETDATA SETPARAMS FEEDBACK
```

The new parameters will take effect the next time the system boots.

14.7 Managing System Parameters with the System Management Utility (SYSMAN)

Note

Digital recommends you use AUTOGEN to modify system parameters. For more information, see Section 14.5. If you want to view system parameters for a group of nodes or change parameters temporarily use the System Management Utility (SYSMAN).

The System Management utility (SYSMAN) provides the ability to inspect and modify system parameters for an entire VMScluster or for any group of nodes, rather than just one system. The PARAMETERS commands available in SYSMAN duplicate the parameter functions of the OpenVMS System Generation utility (SYSGEN).

You can use SYSMAN to manage system parameters as follows:

Task	For More Information
Show parameter values	Section 14.7.2
Modify current values in the parameter file	Section 14.7.3
Modify active values on a running system ¹	Section 14.7.4

¹Applies only to the dynamic system parameters.

SYSMAN provides the commands and functions shown in Table 14-3.

Table 14-3 SYSMAN PARAMETERS Commands

Command	Function
PARAMETERS SHOW	Displays parameter values. Requires the name of the parameter.

(continued on next page)

Table 14–3 (Cont.) SYSMAN PARAMETERS Commands

Command	Function
PARAMETERS USE	Reads a set of parameters from memory or disk into the work area for inspection or modification. Requires a filename or the additional parameters ACTIVE or CURRENT.
PARAMETERS SET	Changes parameter values only in the work area; more permanent modification requires the PARAMETERS WRITE command. Requires the name and value of the parameter.
PARAMETERS WRITE	Writes the content of the work area to memory or to disk. Requires a filename or the additional parameters ACTIVE or CURRENT.

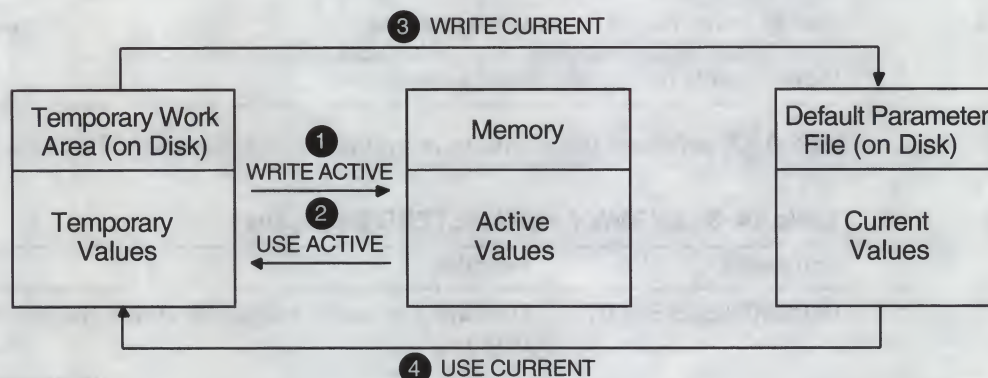
For more information about the temporary work area, see the next section.

14.7.1 Understanding Parameter Values and SYSMAN

It helps to understand the different system parameter values explained in Section 14.1.1. Briefly, **current values** are values stored in the default parameter file on disk. **Active values** are values that are stored in memory and used while the system is running. In addition to these values, SYSMAN writes a temporary copy into its own work area on disk. Figure 14–2 illustrates these different sets of values and how SYSMAN commands affect them. In this figure:

- ❶ WRITE ACTIVE writes temporary parameter values to memory.
- ❷ USE ACTIVE reads values from memory into the work area, where you can modify them.
- ❸ WRITE CURRENT writes temporary parameter values to disk, where they become current values. They become active the next time the system boots.
- ❹ USE CURRENT reads the current values from disk into the work area, where you can modify them.

Figure 14–2 SYSMAN Temporary, Active, and Current Parameter Values



ZK-5275A-GE

14.7 Managing System Parameters with the System Management Utility (SYSMAN)

In a typical session, you can display and change values in the following sequence during a typical session:

1. Read values into SYSMAN's temporary work space with the USE command. USE ACTIVE reads in active values. USE CURRENT reads in current values.
2. Display the parameter values with the SHOW command.
3. Change a value with the SET command. You must use the WRITE command to activate the value.
4. Make the change effective with the WRITE command.
WRITE ACTIVE writes the value to the set of active values. (You can change an active value only if the parameter is a dynamic parameter.) WRITE CURRENT writes the value to the set of current values.

For a list of all the system parameters, see the *OpenVMS System Management Utilities Reference Manual*.

14.7.2 Showing Parameter Values with SYSMAN

You can use the SYSMAN command PARAMETERS SHOW to display parameter values for all the nodes in a cluster.

Examples

1. The following example shows one method to display information about parameters. In this case, using the /LGI qualifier displays all login security control parameters. You can display many categories of parameters, such as /ACP, /ALL, and /SPECIAL. See the *OpenVMS System Management Utilities Reference Manual* for a complete list of parameters and parameter categories.

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS SHOW/LGI
```

Parameter Name	Current	Default	Min.	Max.	Unit	Dynamic
LGI_BRK_TERM	0	1	0	1	Boolean	D
LGI_BRK_DISUSER	0	0	0	1	Boolean	D
LGI_PWD_TMO	30	30	0	255	Seconds	D
LGI_RETRY_LIM	3	3	0	255	Tries	D
LGI_RETRY_TMO	20	20	0	255	Seconds	D
LGI_BRK_LIM	5	5	0	255	Failures	D
LGI_BRK_TMO	300	300	0	-1	Seconds	D
LGI_HID_TIM	300	300	0	-1	Seconds	D

2. This example invokes SYSMAN and specifies the environment to be the local cluster, which consists of NODE21 and NODE22. The example also displays the active value for the LGI_BRK_TMO parameter, which controls the number of seconds that a user, terminal, or node is permitted to attempt login. In this case, it is 600.

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
%SYSMAN-I-ENV, Current command environment:
Clusterwide on local cluster
Username MORIN will be used on nonlocal nodes
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO
```

Parameter Name	Current	Default	Minimum	Maximum	Unit	Dynamic
LGI_BRK_TMO	600	300	0	-1	Seconds	D

Managing System Parameters

14.7 Managing System Parameters with the System Management Utility (SYSMAN)

```
Node NODE22: Parameters in use: ACTIVE
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO         600     300      0        -1 Seconds  D
```

14.7.3 Modifying a Parameter File with SYSMAN

You can use the SYSMAN command **PARAMETERS WRITE** to write system parameter values and the name of the site-independent startup command procedure to your choice of parameter file or the current system parameter file on disk.

The **PARAMETERS WRITE CURRENT** command sends a message to OPCOM to record the event, unless you have changed the system message format with the DCL command **SET MESSAGE**.

Note

The **PARAMETERS WRITE CURRENT** command writes *all* of the active or current parameter values—not just the one you may be working on—to disk.

Examples

1. The following example creates a new parameter specification file:

```
SYSMAN> PARAMETERS WRITE SYS$SYSTEM:NEWPARAM
```

2. When used with the **PARAMETERS SET** command, the **PARAMETERS WRITE** command modifies the current system parameter file on disk:

```
SYSMAN> PARAMETERS SET LGI_BRK_TMO 300
SYSMAN> PARAMETERS WRITE CURRENT
```

14.7.4 Modifying Active Values with SYSMAN

Using the SYSMAN commands **PARAMETERS SET**, **PARAMETERS WRITE**, and **PARAMETERS USE** enables you to modify active parameter values.

Modifying active values immediately affects dynamic parameters by changing their values in memory. The *OpenVMS System Management Utilities Reference Manual* identifies the dynamic parameters, as does the SYSMAN command **PARAMETERS SHOW/DYNAMIC**. Values for nondynamic parameters cannot be changed while the system is running.

Modifying active values does not affect current values in the system parameter file on disk, because the next time you boot the system, the values on disk are established as the active values.

If you set new active parameter values and you want to use the new values for subsequent boot operations, write the new values to the current parameter file with the **PARAMETERS WRITE CURRENT** command, as shown in the Examples section.

Caution

Parameter values modified with SYSMAN will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSMAN, edit the file **SYS\$SYSTEM:MODPARAMS.DAT** as explained in Section 14.5.1 to specify the new parameter values.

14.7 Managing System Parameters with the System Management Utility (SYSMAN)

Managing System Parameters

Examples

1. The following example changes the LGI_BRK_TMO value to 300 in the work area, writes this change into memory as an active value, and displays the active value:

```
SYSMAN> PARAMETERS SET LGI_BRK_TMO 300
```

```
SYSMAN> PARAMETERS WRITE ACTIVE
```

```
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO
```

```
Node NODE21: Parameters in use: ACTIVE
```

```
Parameters in use: ACTIVE
```

Parameter Name	Current	Default	Minimum	Maximum	Unit	Dynamic
LGI_BRK_TMO	300	300	0	-1	Seconds	D

```
Node NODE22: Parameters in use: ACTIVE
```

Parameter Name	Current	Default	Minimum	Maximum	Unit	Dynamic
LGI_BRK_TMO	300	300	0	-1	Seconds	D

2. The following example calls the current parameter values, including LGI_BRK_TMO, from disk to the work area, then displays LGI_BRK_TMO. In this example, the current value on disk is 600.

```
SYSMAN> PARAMETERS USE CURRENT
```

```
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO
```

```
Node NODE21: Parameters in use: CURRENT
```

Parameter Name	Current	Default	Minimum	Maximum	Unit	Dynamic
LGI_BRK_TMO	600	300	0	-1	Seconds	D

```
Node NODE22: Parameters in use: CURRENT
```

Parameter Name	Current	Default	Minimum	Maximum	Unit	Dynamic
LGI_BRK_TMO	600	300	0	-1	Seconds	D

3. The next example writes the LGI_BRK_TMO value of 600 from the work area to memory, where it becomes the active value on the running system. Note that the command **PARAMETER WRITE ACTIVE** writes all the parameter values from the work area into memory, not just the value of LGI_BRK_TMO.

```
SYSMAN> PARAMETERS WRITE ACTIVE
```

```
SYSMAN> PARAMETERS USE ACTIVE
```

```
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO
```

```
Node NODE21: Parameters in use: ACTIVE
```

Parameter Name	Current	Default	Minimum	Maximum	Unit	Dynamic
LGI_BRK_TMO	600	300	0	-1	Seconds	D

```
Node NODE22: Parameters in use: ACTIVE
```

Parameter Name	Current	Default	Minimum	Maximum	Unit	Dynamic
LGI_BRK_TMO	600	300	0	-1	Seconds	D

14.8 Managing System Parameters with the System Generation Utility (SYSGEN)

Note

Digital recommends you use AUTOGEN to modify system parameters. For more information, see Section 14.5. If for some reason you cannot use AUTOGEN, Digital recommends you use the System Management utility (SYSMAN). For more information, see Section 14.7.

Although it is not the recommended method, you can also use the System Generation utility (SYSGEN) to manage system parameters as follows:

Task	For More Information
Show parameter values	Section 14.8.2
Modify current values in the default parameter file	Section 14.8.3
Modify active values on a running system ¹	Section 14.8.4
Create a new parameter file	Section 14.8.5

¹Applies only to the dynamic system parameters.

SYSGEN provides the commands shown in Table 14-4 for managing system parameters. See the SYSGEN section of the *OpenVMS System Management Utilities Reference Manual* for detailed descriptions of SYSGEN commands.

Table 14-4 SYSGEN Commands Used with System Parameters

Command	Function
SHOW	Displays parameter values.
USE	Reads a set of values from memory or disk into a temporary work area for inspection or modification.
SET	Changes parameter values only in the work area; more permanent modification requires the WRITE command.
WRITE	Writes the content of the work area to memory or to disk.

For more information about the temporary work area, see the next section.

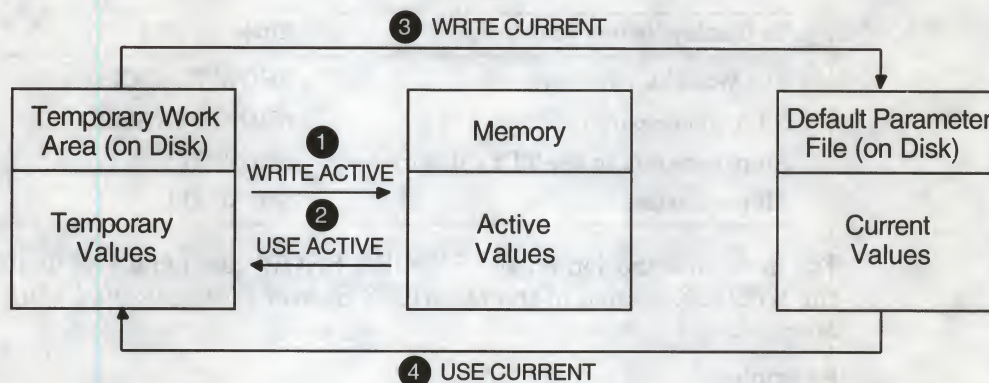
14.8.1 Understanding Parameter Values and SYSGEN

You should understand the different system parameter values explained in Section 14.1.1. Briefly, **current values** are values stored in the default parameter file on disk. **Active values** are values that are stored in memory and used while the system is running. In addition to these values, SYSGEN writes a temporary copy into its own work area on disk. Figure 14-3 illustrates these different sets of values and shows how SYSGEN commands affect them.

Managing System Parameters

14.8 Managing System Parameters with the System Generation Utility (SYSGEN)

Figure 14-3 SYSGEN Temporary, Active, and Current Parameter Values



ZK-5275A-GE

In a typical session, you might display and change values in the following sequence:

1. Read values into SYSGEN's temporary work space with the **USE** command. **USE ACTIVE** reads in active values. **USE CURRENT** reads in current values.
2. Display the parameter values with the **SHOW** command.
3. Change a value with the **SET** command. (Note, however that the **SET** command only changes the value in SYSGEN's temporary work area.)
4. Make the change effective with the **WRITE** command.

WRITE ACTIVE writes the value to the set of active values in memory. (You can change an active value only if the parameter is a dynamic parameter.)
WRITE CURRENT writes the value to the set of current values on disk.

For a list of all the system parameters, see the *OpenVMS System Management Utilities Reference Manual*.

14.8.2 Showing Parameter Values with SYSGEN

To display values for system parameters, perform the following steps:

1. Invoke SYSGEN by entering the following command:
`$ RUN SYS$SYSTEM:SYSGEN`
2. Enter the **USE** command to specify which values you want to display, as follows:

To Display	Enter
Active values	USE ACTIVE
Current values	USE CURRENT
Values from another parameter file	USE file-spec
	For <i>file-spec</i> , specify the parameter file from which you want to display values; for example, USE SYS\$SYSTEM:ALTPARAMS.DAT

3. Enter a **SHOW** command in the following format:

SHOW [/qualifier] [parameter-name]

Managing System Parameters

14.8 Managing System Parameters with the System Generation Utility (SYSGEN)

Specify qualifiers to display parameters grouped by type. For example:

To Display Values For	Enter
The WSMAX parameter	SHOW WSMAX
All dynamic parameters	SHOW/DYNAMIC
All parameters in the TTY category	SHOW/TTY
All parameters	SHOW/ALL

For more information on the SYSGEN SHOW command and qualifiers, see the SYSGEN section of the *OpenVMS System Management Utilities Reference Manual*.

Example

The following example uses SYSGEN to show the current values of all TTY system parameters:

```
$ RUN SYS$SYSTEM:SYSGEN
$ USE CURRENT
SYSGEN> SHOW/TTY
```

```
Parameters in use: Current1
Parameter Name      Current      Default      Min.      Max.      Unit      Dynamic
-----
2
3      4      5      6      7
TTY_SCANDelta      10000000    10000000    100000    -1      100Ns
TTY_DIALTYPE        0           0           0         255     Bit-Encode
TTY_SPEED           15          15          1         16      Special
TTY_RSPEED          0           0           0         16      Special
TTY_PARITY          24          24          0         255     Special
TTY_BUF             80          80          0         65535   Characters
TTY_DEFCHAR         402657952   402657952   0         -1      Bit-Encode
TTY_DEFCHAR2        135178      4098        0         -1      Bit-Encode
TTY_TYPAHDSZ        78          78          0         -1      Bytes
TTY_ALTYPAHD        2048        200         0         32767   Bytes
TTY_ALTALARM        750         64          0         -1      Bytes
TTY_DMASIZE         64          64          0         -1      Bytes      D 8
TTY_CLASSNAME       "TTY"       "TTY"       "AA"      "ZZ"    Ascii
TTY_SILOTIME        8           8           0         255     Ms
TTY_TIMEOUT         3600        900         0         -1      Seconds    D
TTY_AUTOCHAR        7           7           0         255     Character  D
SYSGEN>
```

SYSGEN displays the following information:

- ¹ The values in use (in this example, current values)
- ² The name of the system parameter
- ³ The value requested (in this example, the current value). The heading of this column is always "Current," regardless of whether it displays the current or active value of the parameter. In this context, "Current" refers to the value of this parameter *currently* in use, as specified by the USE command; it does not refer to the *current value* of the parameter stored on disk with the WRITE CURRENT command.
- ⁴ The default value
- ⁵ The minimum value
- ⁶ The maximum value
- ⁷ The unit of allocation

14.8 Managing System Parameters with the System Generation Utility (SYSGEN)

- ⑧ A "D," if the system parameter is dynamic

14.8.3 Modifying the System Parameter File with SYSGEN

Caution

Parameter values modified with the System Generation utility (SYSGEN) will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSGEN, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 14.5.1 to specify the new parameter values.

Note

Although you can modify system parameter values with SYSGEN, Digital recommends you use AUTOGEN. For more information, see Section 14.5.

If you cannot use AUTOGEN, Digital recommends you use the System Management utility (SYSMAN) to modify system parameters. For more information, see Section 14.7.

Modifying the current values in the default system parameter file has no immediate effect on active values on a running system. However, during subsequent boot operations, the system is initialized with the new values.

Example

The following example modifies the TTY_TIMEOUT parameter value in the VAX system parameter file:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SET TTY_TIMEOUT 3600
SYSGEN> WRITE CURRENT
%OPCOM, 15-APR-1995 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITECUR, CURRENT system parameters modified by process
ID 00160030 into file VAXVMSYS.PAR
SYSGEN> EXIT
```

14.8.4 Modifying Active Values with SYSGEN

Caution

Parameter values modified with SYSGEN will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSGEN, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 14.5.1 to specify the new parameter value.

Note

Although you can modify system parameter values with SYSGEN, Digital recommends you use AUTOGEN or the System Management utility (SYSMAN). For more information, see Section 14.7.

Managing System Parameters

14.8 Managing System Parameters with the System Generation Utility (SYSGEN)

Modifying active values immediately affects dynamic parameters by changing their values in memory. The *OpenVMS System Management Utilities Reference Manual* identifies the dynamic parameters (as does the SYSGEN command SHOW/DYNAMIC). Values for nondynamic parameters cannot be changed while the system is running.

Modifying active values does not affect the current values in the system parameter file on disk. The next time you boot the system, the old current values are established as the active values.

If you set new active parameter values (by entering WRITE ACTIVE) and you want to use the new values for subsequent boot operations, you must write the new values to the current parameter file on disk by entering the WRITE CURRENT command, as explained in Section 14.8.3. If the parameters are not dynamic parameters, you must enter the WRITE CURRENT command and reboot the system.

When you change active parameters with SYSGEN, the operator communication manager (OPCOM) writes a message to the operator log and the operator console, unless you have changed the system message format with the DCL command SET MESSAGE.

Examples

1. The following example modifies the active value of the PFCDEFAULT parameter:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> SET PFCDEFAULT 127
SYSGEN> WRITE ACTIVE
%OPCOM, 15-APR-1995 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITEACT, ACTIVE system parameters modified by process
ID 00160030
SYSGEN> EXIT
```

2. The following example modifies the active value of the PFCDEFAULT parameter and also writes it to the Alpha system parameter file, so it will be used when the system reboots:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> SET PFCDEFAULT 127
SYSGEN> WRITE ACTIVE
%OPCOM, 15-APR-1995 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITEACT, ACTIVE system parameters modified by process
ID 00160030
SYSGEN> WRITE CURRENT
%OPCOM, 15-APR-1995 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITECUR, CURRENT system parameters modified by process
ID 00160030 into file ALPHAVMSSYS.PAR
SYSGEN> EXIT
```

14.8.5 Creating a New Parameter File with SYSGEN

Creating a new parameter file has no effect on the running system. During a subsequent conversational boot operation, however, you can initialize the active system with the values of the new file.

Managing System Parameters

14.8 Managing System Parameters with the System Generation Utility (SYSGEN)

How to Perform This Task

1. Invoke SYSGEN by entering the following commands:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
```

2. Enter a command in the following format to write a copy of a parameter file into SYSGEN's temporary workspace:

USE *file-spec*

Where *file-spec* is the file specification for the parameter file to be used as a base. You will modify the values in this file to create a new parameter file.

3. Enter commands in the following form to modify values as needed:

SET *parameter-name* *parameter-value*

For *parameter-name*, specify the name of the parameter to be changed. For *parameter-value*, specify the new value.

4. Specify a command in the following format to write the values to a new parameter file:

WRITE *file-spec*

where *file-spec* is the file specification for the parameter file to be created.

5. Exit SYSGEN.

Caution

Parameter values modified with the System Generation utility (SYSGEN) will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSGEN, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 14.5.1 to specify the new parameter values.

Examples

1. The following example creates a new version of the parameter file PARAMS.PAR with a new value for the TTY_TIMEOUT parameter:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> USE SYS$MANAGER:PARAMS.PAR
SYSGEN> SET TTY_TIMEOUT 3600
SYSGEN> WRITE SYS$MANAGER:PARAMS.PAR
SYSGEN> EXIT
```

2. The following example creates a file named SYS\$SYSTEM:OURSITE.PAR, using the PARAMS.PAR file as a base:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> USE SYS$MANAGER:PARAMS.PAR
SYSGEN> SET TTY_TIMEOUT 1000
SYSGEN> WRITE OURSITE.PAR
SYSGEN> EXIT
```


14.9 Modifying System Parameters with a Conversational Boot

Note

Although you can modify system parameters with a conversational boot, Digital recommends you use AUTOGEN or the System Management utility (SYSMAN). For more information, see Section 14.5 and Section 14.7.

Use a conversational boot only to change isolated system parameters *temporarily* or in an emergency. For example, during a system upgrade, you would use a conversational boot to modify STARTUP_P1 to use a minimum startup.

Remember that if you change a value and do not add the changed value to the AUTOGEN parameter file MODPARAMS.DAT, AUTOGEN will overwrite the value the next time AUTOGEN executes.

With a conversational boot operation, you can modify the active parameter values in the following ways before the system boots:

Task	For More Information
Modify active values for individual parameters	Section 4.2.1
Initialize active values using values stored in a parameter file other than the default parameter file	Section 4.2.2
Reinitialize active values using default values	Section 4.3.1

At the end of the conversational boot, the default system parameter file is modified to store the new active parameter values.

Caution

Parameter values modified with a conversational boot will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with a conversational boot, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 14.5.1 to specify the new parameter values.

Managing System Page, Swap, and Dump Files

The system page, swap, and dump files are created by default. However, you should understand these files. In addition, you might want to change them to meet the needs of your site.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Displaying information about page and swap files	Section 15.3
Calculating appropriate sizes for files	Section 15.4
Minimizing dump file size when disk space is insufficient	Section 15.5
Using SDA to analyze the contents of a crash dump	Section 15.6
‡Using SDA CLUE commands to obtain and analyze summary crash dump information	Section 15.7
†Using CLUE to obtain historical information about crash dumps	Section 15.8
Copying dump files to tape or disk	Section 15.9
Saving the contents of the system dump file after a system failure	Section 15.11
Freeing dump information from the page file	Section 15.12
Creating page and swap files	Section 15.13
Installing page and swap files	Section 15.14
Removing page, swap, and dump files	Section 15.15
Changing page, swap, and dump file sizes	Section 15.16
Controlling page, swap, and dump file sizes in MODPARAMS.DAT	Section 15.16.1.1

†VAX specific

‡Alpha specific

This chapter explains the following concepts:

Concept	Section
Understanding the system dump file	Section 15.1
Understanding page and swap files	Section 15.2
‡Understanding SDA CLUE	Section 15.7.1
†Understanding CLUE	Section 15.8.1

†VAX specific

‡Alpha specific

15.1 Understanding the System Dump File

When the operating system detects an unrecoverable error or an inconsistency within itself that causes the system to fail, it writes the contents of the error log buffers, processor registers, and memory into the **system dump file**, overwriting its previous contents.

When writing the system dump file, the system displays a number of console messages and information about the error or inconsistency. The following message tells you that the dump file was successfully written:

System dump complete

Caution

Be sure to wait until the system dump file is complete and you see this message before using the console terminal to halt the system. If you don't, your system might not save a complete dump file.

The contents of the console messages and the contents of the system dump file are important sources of information in determining the cause of a system failure. You use the contents in the following ways:

- Use the System Dump Analyzer utility (SDA) to analyze the contents of the dump and determine the cause of a failure.
- On Alpha systems, use SDA CLUE commands to obtain and analyze summary dump file information. ♦
- On VAX systems, use CLUE to obtain historical information from system dump files. ♦
- Send the contents of the dump to Digital Equipment Corporation, along with a Software Performance Report (SPR).

Alpha

VAX

The default system dump file, SYS\$SPECIFIC:[SYSEXE]SYSDUMP.DMP, is furnished as an empty file in the operating system distribution kit.

AUTOGEN automatically determines an appropriate size for the system dump file for your hardware configuration and system parameters. Refer to Section 15.5 for information on minimizing dump file size if disk space is insufficient. For special configurations or varying work loads you might want to change the size of the system dump file. For information, see Section 15.16.1.

You do not need a system dump file to run the operating system. However, you must have system dump file to diagnose system crashes.

Using the Page File to Store System Crash Dumps

The operating system uses the latest version of SYS\$SYSTEM:SYSDUMP.DMP to store system crash dumps. If SYSDUMP.DMP does not exist in SYS\$SYSTEM, the operating system uses the system paging file, SYS\$SYSTEM:PAGEFILE.SYS, overwriting the contents of that file. If the SAVEDUMP system parameter is set, the crash dump is retained in PAGEFILE.SYS when the system is booted. If SAVEDUMP is clear, the system uses the paging file for paging and any dump written to the paging file is lost.

Managing System Page, Swap, and Dump Files

15.1 Understanding the System Dump File

If you use SYS\$SYSTEM:PAGEFILE.SYS to capture system crash dumps, you should later free the space occupied by the dump for use in system paging, with either of the following methods:

- Use the SDA COPY command to copy the page file to a different file.
- Use the SDA RELEASE command to delete the information from the page file.

For detailed instructions, see Section 15.12.

Include the appropriate SDA command in the SYSTARTUP_VMS.COM startup command procedure to free dump information from the page file each time the system reboots.

Caution

Be careful when using the page file for selective dumps. Selective dumps use up all available space. If your page file is small, selective dump information might fill the entire page file, leaving no space for paging during system boot. This can cause the system to hang during reboot.

Types of Dumps

The two types of dumps are : physical and selective. Table 15–1 defines physical and selective dumps. Table 15–3 compares the information available in physical and selective dump files.

Table 15–1 Comparison of Physical and Selective Dumps

Type	Description
Physical dump	Writes the entire contents of physical memory to the dump file. To ensure a useful physical dump, the dump file must be large enough to contain all of physical memory.
Selective dump	Stores those portions of memory most likely to be useful in crash dump analysis. A selective dump is useful when disk space is not available to hold all of physical memory. To direct your system to save a selective dump, set the system parameter DUMPSTYLE to the appropriate value. For more information, see Section 15.5 and also the <i>OpenVMS System Management Utilities Reference Manual</i> .

Requirements for Creating a Useful System Dump

The following requirements must be met for the operating system to write a useful system dump file:

- The system parameter DUMPBUG must be set to 1 (the default value).
- If the system parameter SAVEDUMP is set to 0 (the default) the file SYS\$SPECIFIC:[SYSEXE]SYSDUMP.DMP must exist on the system disk.
- If the file SYS\$SPECIFIC:[SYSEXE]SYSDUMP.DMP does not exist on the system disk, the page file must be used to store the dump. The system parameter SAVEDUMP must be set to 1 and the file SYS\$SPECIFIC:[SYSEXE]PAGEFILE.SYS must exist on the system disk.

Managing System Page, Swap, and Dump Files

15.1 Understanding the System Dump File

- If sufficient disk space is not available to allow a system dump file that can hold all of memory, the system parameter DUMPSTYLE must be set to the appropriate value to store a selective dump. For more information, see Section 15.5.
- The system dump file (or page file if the SAVEDUMP system parameter is set) must be large enough to hold all information that is to be written if the system fails.

If the system parameter DUMPSIZE is set, AUTOGEN automatically sizes SYSDUMP.DMP if enough disk space is available.

If the system parameter SAVEDUMP is set, AUTOGEN performs no operations on the dump file.

AUTOGEN sizes the page file only for paging use, regardless of whether the SAVEDUMP system parameter is set.

BACKUP Considerations

System dump files have the NOBACKUP attribute, so the Backup utility (BACKUP) does not copy them unless you use the qualifier /IGNORE=NOBACKUP when invoking BACKUP. When you use the SDA COPY command to copy the system dump file to another file, the operating system does not automatically set the new file to NOBACKUP. If you want to set the NOBACKUP attribute on the copy, use the SET FILE command with the /NOBACKUP qualifier as described in the *OpenVMS DCL Dictionary*.

Security Considerations

As included in the distribution kit, SYS\$SYSTEM:SYSDUMP.DMP is protected against world access. Because a system dump file can contain privileged information, you should keep this level of protection on dump files. Similarly, when you copy dump files using the System Dump Analyzer utility (SDA) as explained in Section 15.11 and Section 15.12, be sure to protect the copy from world read access. For more information on file protection, see the *Security Guide*.

15.2 Understanding Page and Swap Files

As part of memory management, the operating system makes efficient use of physical memory by moving information between physical memory and files stored on disk. The system does this in two ways: **paging** and **swapping**. Table 15–2 defines these and related terms.

Table 15–2 Paging and Swapping Terminology

Term	Definition
Paging	To efficiently use the physical memory allotted to a <i>process</i> , the operating system moves infrequently used portions of a process workspace out of physical memory to a file. For more information on paging, see the <i>Guide to OpenVMS Performance Management</i> .

(continued on next page)

Managing System Page, Swap, and Dump Files

15.2 Understanding Page and Swap Files

Table 15–2 (Cont.) Paging and Swapping Terminology

Term	Definition
Page file	The file to which the system writes paged portions of memory. Your distribution kit includes a page file named <code>SYS\$SYSTEM:PAGEFILE.SYS</code> . If necessary, <code>SYS\$SYSTEM:PAGEFILE.SYS</code> can be used in place of the system crash dump file. For more information, see Section 15.1.
Swapping	To efficiently use the physical memory available for the <i>entire system</i> , the operating system moves the entire workspace of a less active process out of physical memory to a file. For more information on swapping, see the <i>Guide to OpenVMS Performance Management</i> .
Swap file	The file to which the system writes swapped portions of memory. Your distribution kit includes a swap file named <code>SYS\$SYSTEM:SWAPFILE.SYS</code> .
Primary page and swap files	The default page and swap files provided with your distribution kit. These files are named <code>SYS\$SYSTEM:PAGEFILE.SYS</code> and <code>SYS\$SYSTEM:SWAPFILE.SYS</code> .
Secondary page and swap files	Additional page and swap files that you might create for performance or disk space reasons. If you kept the primary page and swap file on the system disk, the system uses the space in the secondary files for paging and swapping in addition to the space in the primary page and swap files. For information on creating secondary page and swap files, see Section 15.13.

Installing Files

Page and swap files must be installed before the system can use them. The system automatically installs the latest versions of `SYS$SYSTEM:PAGEFILE.SYS` and `SWAPFILE.SYS` during startup. If you create secondary page and swap files, you must make sure the system installs them during startup. For more information on installing page and swap files, see Section 15.14.

File Sizes and Locations

AUTOGEN automatically determines appropriate sizes for the files for your hardware configuration and system parameters. For special configurations or varying work loads, you might want to change the size of the page or swap file. For information, see Section 15.16.1.

If your system does not require the page file for storing crash dumps, you can move it off the system disk. However, you should keep one page file on the system disk, if possible, so that you can boot the system if another disk holding the page files becomes unavailable. The swap file can also be moved off the system disk.

Managing System Page, Swap, and Dump Files

15.3 Displaying Information About Page and Swap Files

15.3 Displaying Information About Page and Swap Files

The DCL command `SHOW MEMORY/FILES` displays information about the page and swap files existing on your system, including file names, sizes, and the amount of space used. For example:

```
$ SHOW MEMORY/FILES
System Memory Resources on 12-MAY-1995 11:54:20.06

Paging File Usage (pages):
DISK$PAGE:[SYSEXE]SWAPFILE_IPL31.SYS;2      Free  Reservable  Total
DISK$PAGE:[SYSEXE]PAGEFILE_IPL31.SYS;1      79992    79992    79992
                                           23263   -370027   249992
```

Note that the number displayed in the column labeled "Reservable" can be a negative number. Processes can reserve more space than is available because it is unlikely that all the reserved space will be used for paging at one time.

15.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

When you install or upgrade the operating system, AUTOGEN automatically calculates appropriate sizes for your system page, swap, and dump files based on your hardware configuration and system parameters. However, you might want to manually calculate the sizes for these files. The following sections describe how to determine appropriate sizes for the system page, swap, and dump files.

VAX

If you are running from a saved Snapshot image, changing the size of any of the page, swap, or dump files disables the ability to boot from that image, and you must create a new Snapshot image. For more information, see Section 4.7.♦

15.4.1 Calculating System Dump File Size

Sufficient space in the system dump file is critical to saving a complete crash dump. The AUTOGEN command procedure calculates an appropriate size for your dump file. However, if you want to manually calculate the dump file size, use the following formula, which calculates the file size required to hold a physical dump.

For SYSDUMP.DMP

VAX

On VAX systems, use the following formula:

```
size-in-blocks(SYS$SYSTEM:SYSDUMP.DMP)
= size-in-pages(physical-memory)
+ (number-of-error-log-buffers * blocks-per-buffer)
+ 1 ♦
```

Alpha

On Alpha systems, use the following formula:

```
size-in-blocks(SYS$SYSTEM:SYSDUMP.DMP)
= size-in-pages(physical-memory)
* blocks-per-page
+ (number-of-error-log-buffers * blocks-per-buffer)
+ 2 ♦
```

where:

size-in-pages

Is the size of physical memory, in pages. Use the DCL command `SHOW MEMORY` to determine the total size of physical memory on your system.

Managing System Page, Swap, and Dump Files

15.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

blocks-per-page

Is the number of blocks per page of memory.

On VAX systems, the disk block size and page size are identical (512).

On Alpha systems, calculate the number of blocks per page of memory by dividing the system's page size by 512 (the size of a block). Use the following commands:

```
$ PAGESIZE==F$GETSYI ("PAGE_SIZE")
$ BLOCKSPERPAGE=PAGESIZE/512
$ SHOW SYMBOL BLOCKSPERPAGE
```

number-of-error-log-buffers

Is the value of the system parameter **ERRORLOGBUFFERS**. This parameter sets the number of error log buffers to permanently allocate in memory.

blocks-per-buffer

Is the value of the system parameter **ERLBUFFERPAGES**. This parameter sets the number of pages of memory in each buffer.

A large memory system or a system with small disk capacity may not be able to supply enough disk space for a full memory dump. Under these circumstances, you should set the system parameter **DUMPSTYLE** to the appropriate value to indicate that the system is to dump only selective information. For more information, see Section 15.5.

For PAGEFILE.SYS

If **SYS\$SYSTEM:SYSDUMP.DMP** does not exist, the system writes crash dumps to the primary page file **SYS\$SYSTEM:PAGEFILE.SYS**. The **AUTOGEN** command procedure calculates an appropriate size for your page file. However, if you want to manually calculate the minimum page file size required to hold crash dumps, use the following formula:

VAX

On VAX systems:

```
size-in-blocks(SYS$SYSTEM:PAGEFILE.SYS)
= size-in-pages(physical-memory)
+ (number-of-error-log-buffers * blocks-per-buffer)
+ 1
+ 1000 ♦
```

Alpha

On Alpha systems:

```
Size-in-blocks(SYS$SYSTEM:PAGEFILE.SYS)
= size-in-pages(physical-memory)
* blocks-per-page
+ (number-of-error-log-buffers * blocks-per-buffer)
+ 2
+ RSRVPAGCNT ♦
```

where:

size-in-pages

Is the size of physical memory, in pages. Use the DCL command **SHOW MEMORY** to determine the total size of physical memory on your system.

Managing System Page, Swap, and Dump Files

15.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

blocks-per-page

Is the number of blocks per page of memory.

On VAX systems, the disk block size and page size are identical (512).

On Alpha systems, calculate the number of blocks per page of memory by dividing the system's page size by 512 (the size of a block). Use the following commands:

```
$ PAGESIZE==F$GETSYI ("PAGE_SIZE")
$ BLOCKSPERPAGE=PAGESIZE/512
$ SHOW SYMBOL BLOCKSPERPAGE
```

number-of-error-log-buffers

Is the value of the system parameter `ERRORLOGBUFFERS`. This parameter sets the number of error log buffers to permanently allocate in memory.

blocks-per-buffer

Is the value of the system parameter `ERLBUFFERPAGES`. This parameter sets the number of pages of memory in each buffer.

`RSRVPAGCNT`

Is the value of the `RSRVPAGCNT` special system parameter.

Caution

This formula calculates only the minimum size requirement for saving a dump in the system's primary page file. For most systems, the page file must be larger than this to avoid hanging the system. For more information about calculating the page file size, see Section 15.4.2.

15.4.2 Calculating Page File Size

Sufficient page file space is critical to system performance. The `AUTOGEN` command procedure calculates an appropriate size for your page file space. The size calculated by `AUTOGEN` should be sufficient. However, if you want to manually calculate the size for page file space, use one of the following formulas.

VAX

On VAX systems:

```
size-in-blocks (total for all page files on the system)
= size-of-average-process (in pages)
* maximum-number-of-processes ♦
```

Alpha

On Alpha systems:

```
size-in-blocks (total for all page files on the system)
= size-of-average-process (in pagelets)
* maximum-number-of-processes ♦
```

where:

size-of-average-process

Is the value of the average virtual size of the process. Use the following command to find the virtual size of the process.

```
$ SHOW PROCESS/CONTINUOUS/ID=process-id
```

On VAX systems, specify this value in blocks.

On Alpha systems, specify this value in pagelets.

Managing System Page, Swap, and Dump Files

15.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

maximum-number-of-processes Is the value of the MAXPROCESSCNT system parameter.

The size you calculate can be represented in one of the following ways:

- In the primary page file only
- Distributed across primary and secondary page files
- If you have removed the primary page file in SYS\$SYSTEM, distributed across a number of secondary page files

Once you have determined an initial size for your page file or files (either with AUTOGEN, or manually), monitor page file usage by executing AUTOGEN with the following command:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS TESTFILES FEEDBACK
```

With this command, AUTOGEN writes page file usage and size recommendations to the feedback report AGEN\$PARAMS.REPORT. (For more information on AUTOGEN and the feedback report, see Section 14.4 and Section 14.4.2.) The DCL command SHOW MEMORY/FILES also displays file usage, as explained in Section 15.2.

Keep page file usage less than half the size of the page file or files. If a paging file starts to fill to the point where system performance is being affected, a message will be printed on the console terminal. If this happens, increase the size of your page file or files or install additional files.

Note

Your system resources and work load affect the required size of your page file. You should be familiar with your system resources and work load. For more information, see the *Guide to OpenVMS Performance Management*.

You limit the amount of page file space consumed by user programs by using the /PGFLQUOTA qualifier of the AUTHORIZE commands ADD and MODIFY. (See the AUTHORIZE section in the *OpenVMS System Management Utilities Reference Manual* for more information.) Do not reduce the value of /PGFLQUOTA below 1024. Size requirements of the page file vary widely, depending on user applications.

15.4.3 Calculating Swap File Size

Sufficient swap file space is critical to system performance. The AUTOGEN command procedure calculates an appropriate size for your swap file space. If you want to manually calculate the size for swap file space, use the following formula:

```
size-in-blocks (total for all swap files on the system)
= maximum-number-of-processes
* average-working-set-quota-of-processes-on-system
```

where:

maximum-number-of-processes Is the value of the MAXPROCESSCNT system parameter.

Managing System Page, Swap, and Dump Files

15.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

average-working-set-quota-of-processes-on-system

Is the average value of the WSQUOTA limit for processes running on the system.

On VAX systems, specify the value in pages.

On Alpha systems, specify the value in pagelets.

On Alpha and VAX systems, the size you calculate can be represented in any of the following ways:

- In the primary swap file only
- Distributed across primary and secondary swap files
- If you have removed the primary swap file in SYS\$SYSTEM, distributed across a number of secondary swap files

Once you have determined an appropriate size for swapfile space (either manually or with AUTOGEN), monitor swap file usage with the DCL command SHOW MEMORY/FILES as explained in Section 15.3. Keep at least one-third of the swap file space unused; otherwise, system performance can be severely affected.

Note

Your system resources and work load determine the required size of your swap file. You should be familiar with your system resources and work load. For more information, see the *Guide to OpenVMS Performance Management*.

15.5 Minimizing Dump File Size When Disk Space Is Insufficient

In certain system configurations, it may be impossible to preserve the entire contents of memory in a disk file. For instance, a large memory system may not be able to supply enough disk space for a full memory dump. If your system attempts to save all of memory but the dump file is too small to accommodate the entire dump, the System Dump Analyzer utility (SDA) might not be able to analyze the dump.

VAX

On VAX systems, insufficient dump space would also prevent the Crash Logger Utility Extractor (CLUE) from being able to analyze the dump. ♦

On VAX and Alpha systems, to preserve those portions of memory that contain information most useful in determining the causes of system failures, you can use selective dumps. Table 15–1 defines physical and selective dumps. Table 15–3 compares the information available in physical and selective dump files.

Table 15–3 Comparison of Physical and Selective Dump Files

Type	Available Information	Unavailable Information
Physical dump	Complete contents of physical memory in use, stored in order of increasing physical address and error log buffers.	Contents of paged-out memory at the time of the crash.

(continued on next page)

Managing System Page, Swap, and Dump Files

15.5 Minimizing Dump File Size When Disk Space Is Insufficient

Table 15–3 (Cont.) Comparison of Physical and Selective Dump Files

Type	Available Information	Unavailable Information
Selective dump	System page table, global page table, system space memory, error log buffers, and process and control regions (plus global pages) for all saved processes.	Contents of paged-out memory at the time of the crash, process and control regions of unsaved processes, and memory not mapped by a page table.

To direct your system to save selective dumps, set the system parameter `DUMPSTYLE` to the appropriate value. System parameters and their values are in the appendix of the *OpenVMS System Management Utilities Reference Manual*. For information on how to change system parameter values, see Section 14.5.

15.6 Using SDA to Analyze the Contents of a Crash Dump

The System Dump Analyzer utility (SDA) lets you interpret the contents of the dump file to investigate the probable causes of the crash. For information on analyzing a crash dump, see the *System Dump Analyzer Utility Manual*.

If your system fails, you should send Digital Equipment Corporation a Software Performance Report (SPR) and a copy of the system dump file written at the time of the failure. For information on copying the system dump file, see Section 15.9.

15.7 Using SDA CLUE Commands to Analyze Crash Dump Files (Alpha Only)

Alpha

SDA CLUE (Crash Log Utility Extractor) commands automate the analysis of crash dumps and maintain a history of all fatal bugchecks on a standalone system or cluster. SDA CLUE commands can be used in conjunction with SDA to collect and decode additional dump file information not readily accessible through standard SDA.

15.7.1 Understanding CLUE (Alpha Only)

On Alpha systems, SDA is automatically invoked by default when you reboot the system after a system failure. To better facilitate crash dump analysis, SDA CLUE commands automatically capture and archive summary dump file information in a CLUE listing file.

A startup command procedure initiates commands that:

- Invoke SDA
- Issue an SDA CLUE HISTORY command
- Create a listing file called `CLUE$nodename_ddmmyy_hhmm.LIS`

The CLUE HISTORY command adds a one-line summary entry to a history file and saves the following output from SDA CLUE commands in the listing file:

- Crash dump summary information
- System configuration
- Stack decoder
- Page and swap files
- Memory management statistics

Managing System Page, Swap, and Dump Files

15.7 Using SDA CLUE Commands to Analyze Crash Dump Files (Alpha Only)

- Process DCL recall buffer
- Active XQP processes
- XQP cache header

The contents of this CLUE list file can help you analyze a system failure.

If these files accumulate more space than the threshold allows (default 5000 blocks), the oldest files are deleted until the threshold limit is reached. This can also be customized using the CLUE\$MAX_BLOCK logical name.

To inhibit the running of CLUE at system startup, define the logical CLUE\$INHIBIT in the SYLOGICALS.COM file as /SYS TRUE.

It is important to remember that CLUE\$nodename_ddmmyy_hhmm.LIS contains only an overview of the crash dump and does not always contain enough information to determine the cause of the crash. If you must do an indepth analysis of the system crash, Digital recommends that you always use the SDA COPY command to save the dump file.

15.7.2 Displaying Data Using SDA CLUE Commands (Alpha Only)

Invoke CLUE commands at the SDA prompt as follows:

```
SDA> CLUE CONFIG
```

CLUE commands provide summary information of a crash dump captured from a dump file. When debugging a crash dump interactively, you can use SDA CLUE commands to collect and decode some additional information from a dump file, which is not easily accessible through standard SDA. For example, CLUE can quickly provide detailed XQP summaries.

You can also use CLUE commands interactively on a running system to help identify performance problems.

You can use all CLUE commands when analyzing crash dumps; the only CLUE commands that are not allowed when analyzing a running system are CLUE CRASH, CLUE ERRLOG, CLUE HISTORY, and CLUE STACK. ♦

15.8 Using CLUE to Obtain Historical Information About Crash Dumps (VAX Only)

VAX

On VAX systems, the Crash Log Utility Extractor (CLUE) is a tool for displaying the contents of a **crash history file**. By examining the contents of the crash history file, you can understand and resolve the issues responsible for failures (crashes), and you might also obtain other useful data.

15.8.1 Understanding CLUE (VAX Only)

The crash history file, which is created and updated by CLUE, contains key parameters from crash dump files. Unlike crash dumps, which are overwritten with each system failure and are therefore typically available only for the most recent failure, the crash history file is a permanent record of system failures.

After a system fails and physical memory is copied to the crash dump file, CLUE automatically appends the relevant parameters to the file CLUE\$OUTPUT:CLUE\$HISTORY.DATA when the system is restarted. The remainder of this section describes how you can use CLUE to display the data it has collected; reference information about CLUE is available in the *OpenVMS System Management Utilities Reference Manual*.

15.8 Using CLUE to Obtain Historical Information About Crash Dumps (VAX Only)

Note

The history file will typically grow by about 10-15 blocks for each entry. You can limit the number of entries in the binary file by defining the logical name CLUE\$MAX_ENTRIES to be the maximum number desired. When this number is reached, the oldest entries are deleted from the history file.

By default, operator shutdowns are recorded in the history file. You can exclude information from operator shutdowns in the history file by defining the logical name CLUE\$EXCLUDE_OPERS as being TRUE, for example by including the following line in SYS\$MANAGER:SYSTARTUP_VMS.COM:

```
$ DEFINE /SYSTEM CLUE$EXCLUDE_OPERS TRUE
```

15.8.2 Displaying Data Using CLUE (VAX Only)

To display data using CLUE, you must first define the following symbol:

```
$ CLUE ::= $CLUE
```

After defining the symbol, you can use CLUE to display information by entering the following command:

```
$ CLUE/DISPLAY
CLUE_DISPLAY>
```

At the CLUE_DISPLAY> prompt, you can issue commands to do the following:

- Use the DIRECTORY command to list failures that have occurred since a specified date, failures of a particular type, failures that contain a specified module, and failures that have a specified offset.

For example, you can list all the failures in the history file using the DIRECTORY command, as follows:

```
CLUE_DISPLAY> DIRECTORY
```

- Use the SHOW command to generate information similar to that obtained from certain commands in the System Dump Analyzer utility (SDA).

For example, if you wanted complete information on the crash listed as crash number 7, the following SHOW command would provide the information:

```
CLUE_DISPLAY> SHOW ALL 7
```

- Use the EXTRACT command to write the data from an entry to a file.

For example, the following command writes the data from entry number 7 in the crash history file to a file named 15MAYCRASH.TXT:

```
CLUE_DISPLAY> EXTRACT 7/OUTPUT=15MAYCRASH.TXT
```

For more information about CLUE commands, see the *OpenVMS System Management Utilities Reference Manual*. ♦

15.9 Copying Dump Files to Tape or Disk

If your system fails, you should send a copy of the contents of the system dump file to Digital Equipment Corporation along with a Software Performance Report (SPR). You can use the Backup utility (BACKUP) to create save sets containing system dump files on magnetic tape or disk. However, when using BACKUP to copy dump files, you must specify the /IGNORE=(NOBACKUP,INTERLOCK) qualifier for the following reasons:

- By default, the system dump file has the NOBACKUP attribute, so it is not copied unless you specify /IGNORE=NOBACKUP.
- The system keeps an open channel to the dump file, so the file is not copied unless you specify /IGNORE=INTERLOCK.

For more information on using BACKUP, see Section 10.13.2. For information on BACKUP commands, see the BACKUP section in the *OpenVMS System Management Utilities Reference Manual*.

15.10 Dump File Off the System Disk (VAX Systems Only)

VAX

You can place the system dump file on a device other than the system disk.

15.10.1 Requirements

Configuring and using this device for writing the system crash dump file are possible based on the following requirements:

- The system must be connected directly to and must boot from CI controllers.
- The dump device must physically connect to the same two HSx CI controllers as the boot device.
- The dump device cannot be MSCP unit zero (0) and only units 1 to 4095 (1—FFF) are supported.

The dump device can be designated on these configurations by using bits 16 through 27 of register 3 (R3). This register can specify the boot device and the desired dump device.

- The dump device directory structure must resemble the current system disk structure. The [SYSn.SYSEXE]SYSDUMP.DMP file will reside there, using the same boot time system root.

You can use AUTOGEN to create this file. In the MODPARAMS.DAT file, the following symbol will prompt AUTOGEN to create the file:

```
DUMPFILe_DEVICE = $nnn$ddcnnnn
```

- The volume label must contain DOSD_DUMP as the first nine characters of the volume label.
- The dump device cannot be part of a volume set.

To enable the bugcheck code and use this file, the DUMPSTYLE system parameter must be correctly enabled. See the *OpenVMS System Management Utilities Reference Manual* for the values. ♦

15.11 Saving the Contents of the System Dump File After a System Failure

15.11 Saving the Contents of the System Dump File After a System Failure

If the system fails, it overwrites the contents of the system crash dump file and the previous contents are lost. For this reason, you should ensure that your system automatically analyzes and copies the contents of the dump file each time the system reboots.

Alpha

On Alpha systems, SDA is invoked by default and a CLUE list file is created. Generated by a set sequence of commands, the CLUE list file contains only an overview of the crash and might not provide enough information to determine the cause of the crash. Digital, therefore, recommends that you always copy the dump file. Please refer to Section 15.7.2 for information on modifying your site-specific command procedure to execute additional commands such as SDA COPY upon startup after a system failure. ♦

VAX

On VAX systems, modify the site-specific startup command procedure SYSTARTUP_VMS.COM so that it invokes the System Dump Analyzer utility (SDA) when the system is booted.

Be aware of the following information:

- When invoked from the site-specific startup procedure in the STARTUP process, SDA executes the specified commands only if the system is booting immediately after a system failure. If the system is rebooting after it was shut down with SHUTDOWN.COM or OPCCRASH.EXE, SDA exits without executing the commands.
- You can use the DCL COPY command to copy the dump file; however, the SDA COPY command is preferred because it marks the dump file as copied. This is particularly important if the dump was written into the paging file, SYS\$SYSTEM:PAGEFILE.SYS, because it releases those pages occupied by the dump to the pager. For more information, see Section 15.12.
- Because a system dump file can contain privileged information, you should protect copies of dump files from world read access. For more information on file protection, see the *Security Guide*.
- System dump files have the NOBACKUP attribute, so the Backup utility (BACKUP) does not copy them unless you use the qualifier /IGNORE=NOBACKUP when invoking BACKUP. When you use the SDA COPY command to copy the system dumpfile to another file, the operating system does not automatically set the new file to NOBACKUP. If you want to set the NOBACKUP attribute on the copy, use the SET FILE command with the /NOBACKUP qualifier as described in the *OpenVMS DCL Dictionary*.

Managing System Page, Swap, and Dump Files

15.11 Saving the Contents of the System Dump File After a System Failure

Example

The SDA COPY command in the following example saves the contents of the file SYS\$SYSTEM:SYSDUMP.DMP and performs some analysis of the file:

```
$ !
$ !      Print dump listing if system just failed
$ !
$ ANALYZE/CRASH_DUMP SYS$SYSTEM:SYSDUMP.DMP
  COPY SYS$SYSTEM:SAVEDUMP.DMP      ! Save dump file
  SET OUTPUT DISK1:SYSDUMP.LIS      ! Create listing file
  READ/EXECUTIVE                    ! Read in symbols for kernel
  SHOW CRASH                        ! Display crash information
  SHOW STACK                        ! Show current stack
  SHOW SUMMARY                      ! List all active processes
  SHOW PROCESS/PCB/PHD/REG          ! Display current process
  EXIT
$ SET FILE/NOBACKUP SYS$SYSTEM:SAVEDUMP.DMP ◆
```

15.12 Freeing Dump Information from the Page File

If you use SYS\$SYSTEM:PAGEFILE.SYS to store a system crash dump, you must later free the space occupied by the dump for use by the pager. If you do not, your system may hang because it has insufficient paging space.

Section 15.1 explains when you might use the page file to store a system crash dump.

How to Perform This Task

1. Invoke the System Dump Analyzer utility (SDA), specifying PAGEFILE.SYS as the target:

```
$ ANALYZE/CRASH_DUMP SYS$SYSTEM:PAGEFILE.SYS
```

2. Enter the SDA command COPY in the following format to copy the dump from SYS\$SYSTEM:PAGEFILE.SYS to another file:

```
COPY dump_filespec
```

For example, to copy the dump file off the system disk to a file called SAVEDUMP.DMP on DISK\$USER5, enter the following command:

```
SDA> COPY DISK$USER5:[DUMPS]SAVEDUMP.DMP
```

Because a system dump file can contain privileged information, you should protect copies of dump files from world read access.

To prevent the system from backing up the complete contents of the file, assign the NOBACKUP attribute to the file with the DCL command SET FILE/NOBACKUP.

Alternatively, to free the pages in the page file that are taken up by the dump without having to copy the dump elsewhere, issue the ANALYZE/CRASH_DUMP/RELEASE command. This command immediately releases the pages to be used for system paging, effectively deleting the dump. Note that this command does *not* allow you to analyze the dump before deleting it.

3. Enter the EXIT command to exit SDA.
4. Include the SDA commands entered in steps 1 and 2 in the site-specific startup command procedure SYSTARTUP_VMS.COM to free page space each time the system reboots.

Managing System Page, Swap, and Dump Files

15.12 Freeing Dump Information from the Page File

Although the DCL COPY command can also be used to copy a dump file, only the SDA COPY command causes the pages occupied by the dump in the system's page file to be released for paging.

Example

The following commands, added to the SYSTARTUP_VMS.COM command procedure, copy the contents of the page file to a file named SAVEDUMP.DMP:

```
$ ANALYZE/CRASH_DUMP SYS$SYSTEM:PAGEFILE.SYS
  COPY DISK$USER5:[DUMPS] SAVEDUMP.DMP
  EXIT
$ SET FILE/NOBACKUP SYS$SYSTEM:SAVEDUMP.DMP
```

15.13 Creating Page and Swap Files

Primary page and swap files are provided in your distribution kit in the following locations:

```
SYS$SYSTEM:PAGEFILE.SYS
SYS$SYSTEM:SWAPFILE.SYS
```

For performance or disk space reasons, you might want to create page and swap files on disks other than the system disk. The following sections explain how to create page and swap files using different methods:

Method	For More Information
Using AUTOGEN (the recommended method)	Section 15.13.1
Using SYSGEN	Section 15.13.2

15.13.1 Using AUTOGEN (Recommended Method)

You can direct AUTOGEN to create new page and swap files by adding symbols to MODPARAMS.DAT to specify the name, location, and size of new files to be created, and running AUTOGEN. Before performing this task, you should understand AUTOGEN and its parameter file MODPARAMS.DAT. For more information, see Section 14.4, which also provides suggested recommendations about when to use AUTOGEN. See also Section 14.4.4.

You can also define symbols in MODPARAMS.DAT to control the size of page, swap, and dump files. For more information, see Section 15.16.1.

How to Perform This Task

1. Add the following symbols to MODPARAMS.DAT to specify the names and locations of the page and swap files to be created:

Definition	For Page Files	For Swap Files
File name and location	PAGEFILE n _NAME = <i>file-spec</i>	SWAPFILE n _NAME = <i>file-spec</i>

For n , use an integer that specifies the page or swap file. Refer to the primary page and swap files by specifying a value of 1 for n ; refer to subsequent files by specifying increasingly higher integer values for n . For example, to refer to a secondary page or swap file, specify a value of 2 for n .

For *file-spec*, specify the full file specification of the file to be created.

Managing System Page, Swap, and Dump Files

15.13 Creating Page and Swap Files

2. Enter the following command to invoke a first pass of AUTOGEN. In this pass, AUTOGEN displays its calculations for system file sizes to SYS\$OUTPUT:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS TESTFILES
```

3. If the file sizes displayed in step 2 are inadequate, add the following symbols to MODPARAMS.DAT to control the size of the files, and return to step 2:

Definition	For Page Files	For Swap Files
File size	MIN_PAGEFILE n _SIZE = <i>block-size</i>	MIN_SWAPFILE n _SIZE = <i>block-size</i>

For n , specify an integer that indicates the page or swap file. Refer to the primary page and swap files by specifying a value of 1 for n ; refer to subsequent files by specifying increasingly higher integer values for n . For example, to refer to a secondary page or swap file, specify a value of 2 for n .

For *block-size*, specify the size in blocks.

4. When you are satisfied with the file sizes displayed in step 2, execute a second pass of AUTOGEN using the following command to install the modified system files when the system is rebooted:

```
$ @SYS$UPDATE:AUTOGEN GENPARAMS REBOOT
```

5. Add commands to the site-specific startup command procedure SYPAGSWPFILES.COM to make sure the files are installed each time the system boots. For instructions, see Section 15.14.

Example

To direct AUTOGEN to create a new secondary swap file named PAGED\$:[PAGESWAP]SWAPFILE.SYS that holds 30,000 blocks, add the following symbols to MODPARAMS.DAT:

```
MIN_SWAPFILE2_NAME = "PAGED$:[PAGESWAP]SWAPFILE.SYS"  
MIN_SWAPFILE2_SIZE = 30000
```

15.13.2 Using SYSGEN

AUTOGEN is the recommended method for creating page and swap files. However, in an emergency, you can use the System Generation utility (SYSGEN) to directly create files. For example, if you see that page file space is becoming dangerously low, you might use SYSGEN to quickly add page file space to prevent the system from hanging.

How to Perform This Task

1. Determine the names, locations, and sizes of the files you plan to create. For information on determining appropriate sizes, see Section 15.4.
2. Invoke SYSGEN by entering the following command:

```
$ RUN SYS$SYSTEM:SYSGEN
```

3. Enter the SYSGEN command CREATE in the following format:

```
CREATE file-spec/SIZE=block-size
```


Managing System Page, Swap, and Dump Files

15.13 Creating Page and Swap Files

For example:

```
SYSGEN> CREATE DUA2:[PAGE_SWAP]PAGEFILE_1.SYS/SIZE=100000
SYSGEN> CREATE DUA2:[PAGE_SWAP]SWAPFILE_1.SYS/SIZE=100000
```

If the file you specify as *file-spec* does not exist, this command creates a file by that name that can be used as a page or swap file. If the file does exist, the command does one of the following:

- If the size you specify is larger than the existing file, the command extends the file.
- If the size you specify is smaller, the command creates a new, smaller file.

For more information on the SYSGEN command CREATE, see the SYSGEN section in the *OpenVMS System Management Utilities Reference Manual*.

4. Install the files, following the instructions in Section 15.14. The system automatically installs the primary page and swap files located in SYS\$SYSTEM. However, other page files are not automatically installed.
5. Add commands to SYS\$MANAGER:SYSPAGSWPFILES.COM to install the files each time the system boots. Follow the instructions in Section 15.14.2.
6. If you do not want AUTOGEN to resize the files according to its calculations, edit MODPARAMS.DAT to specify the sizes of these files. Follow the instructions in Section 15.16.1.1.

Example

The following example uses SYSGEN to create page and swap files. It also installs the files as explained in Section 15.14.

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> CREATE DUA2:[PAGE_SWAP]PAGEFILE_1.SYS/SIZE=100000
SYSGEN> CREATE DUA2:[PAGE_SWAP]SWAPFILE_1.SYS/SIZE=100000
SYSGEN> INSTALL DUA2:[PAGE_SWAP]PAGEFILE_1.SYS/PAGEFILE
SYSGEN> INSTALL DUA2:[PAGE_SWAP]SWAPFILE_1.SYS/SWAPFILE
```

15.14 Installing Page and Swap Files

The system automatically installs the primary page and swap files located in SYS\$SYSTEM. However, other page and swap files are not automatically installed. For this reason, if you create secondary page and swap files, you must also install them with the System Generation utility (SYSGEN). Note that SYSGEN INSTALL commands perform a different function than Install utility (INSTALL) commands.

15.14.1 Installing Interactively

1. Invoke SYSGEN by entering the following command:

```
$ RUN SYS$SYSTEM:SYSGEN
```

2. Enter the SYSGEN command INSTALL in the following format:

```
INSTALL file-spec/filetype
```

For example:

```
SYSGEN> INSTALL DUA2:[PAGE_SWAP]PAGEFILE_1.SYS/PAGEFILE
SYSGEN> INSTALL DUA2:[PAGE_SWAP]SWAPFILE_1.SYS/SWAPFILE
```


Managing System Page, Swap, and Dump Files

15.14 Installing Page and Swap Files

3. To make sure the files are installed each time the system boots, edit `SYS$MANAGER:SYPAAGSWPFILES.COM` to add the commands you entered in step 2. For more information, see Section 15.14.2.

Example

The following example installs page and swap files interactively:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> INSTALL DUA2:[PAGE_SWAP]PAGEFILE_1.SYS/PAGEFILE
SYSGEN> INSTALL DUA2:[PAGE_SWAP]SWAPFILE_1.SYS/SWAPFILE
```

15.14.2 Installing in `SYPAAGSWPFILES.COM`

Page and swap files other than `SYS$SYSTEM:PAGEFILE.SYS` and `SYS$SYSTEM:SWAPFILE.SYS` must be reinstalled each time the system boots. You can do this by adding the commands to install the files to the startup command procedure `SYS$MANAGER:SYPAAGSWPFILES.COM`. The template file `SYS$MANAGER:SYPAAGSWPFILES.TEMPLATE` includes comments that help explain how this file is used.

Before performing this task, you must have created the secondary files, as explained in Section 15.13.

For more information on `SYPAAGSWPFILES.COM`, see Section 5.2.3.

You can also use `SATELLITE_PAGE.COM` to install page and swap files on a VAXcluster or VMScluster satellite node's local disk. `SATELLITE_PAGE.COM` is created when you run `CLUSTER_CONFIG.COM`. For more information on installing page and swap files on a satellite node's local disk, see the *VMScluster Systems for OpenVMS* manual.

How to Perform This Task

1. Invoke any editor to edit `SYS$MANAGER:SYPAAGSWPFILES.COM`.
2. If necessary, add a `MOUNT` command for each disk that holds a page or swap file. This is necessary because only the system disk is mounted at the time `SYPAAGSWPFILES.COM` is invoked.

For example:

```
$ MOUNT/SYSTEM/NOASSIST DUA2: DISK_SYS2
```

For information on the `MOUNT` command, see the *OpenVMS DCL Dictionary*.

The following commands, inserted before the `MOUNT` command, are also useful to determine if the disk is available before mounting. Note, however, that if the disk is broken and cannot mount, these commands will cause an infinite loop.

```
$ LOOP1:
$ ON WARNING THEN GOTO LOOP1
$ WAIT 0000 00:00:00.50
$ READY = F$GETDVI("device:", "AVL")
$ IF READY .EQS. "FALSE" THEN GOTO LOOP1
```

For *device:*, specify the device name.

3. Add the following command to invoke `SYSGEN`:

```
$ RUN SYS$SYSTEM:SYSGEN
```

4. Add commands in the following format to `SYPAAGSWPFILES.COM` to install the files each time the system boots.

Managing System Page, Swap, and Dump Files

15.14 Installing Page and Swap Files

For page files, use the following format:

INSTALL file-spec/PAGEFILE

For example:

INSTALL DUA2:[SYSTEM]PAGEFILE_1.SYS/PAGEFILE

For swap files, use the following format:

INSTALL file-spec/SWAPFILE

For example:

INSTALL DUA2:[SYSTEM]SWAPFILE_1.SYS/SWAPFILE

5. Add an EXIT command to exit SYSGEN:

EXIT

Example

The following example shows commands you might add to SYPAGSWPFILES.COM to install page and swap files named PAGEFILE_1.SYS and SWAPFILE_1.SYS located on the DUA2: device:

```
$ EDIT SYS$MANAGER:SYPAGSWPFILES.COM
[add the following commands to SYPAGSWPFILES.COM:]
.
.
.
$ MOUNT/SYSTEM/NOASSIST DUA2: DISK_SYS2
$ RUN SYS$SYSTEM:SYSGEN
INSTALL DUA2:[SYSTEM]PAGEFILE_1.SYS /PAGEFILE
INSTALL DUA2:[SYSTEM]SWAPFILE_1.SYS /SWAPFILE
EXIT
```

15.15 Removing Page, Swap, and Dump Files

Caution

If you remove a page, swap, or dump file, do not simply delete the file.

How to Perform This Task

1. Use the RENAME command to rename the file to be deleted.
2. Shut down and reboot the system.
3. Delete the file.
4. When you delete a file, make sure you remove from SYPAGESWPFILES.COM and MODPARAMS.DAT any command lines related to the file.

Managing System Page, Swap, and Dump Files

15.15 Removing Page, Swap, and Dump Files

Example

```
$ RENAME DUA2:[SYSTEM]PAGEFILE_1.SYS; DUA2:[SYSTEM]JUNK.SYS;  
$ @SYS$SYSTEM:SHUTDOWN.COM  
.  
.  
.  
[SHUTDOWN.COM shuts down and reboots the system]  
[When the system reboots, log in]  
.  
.  
.  
$ DELETE DUA2:[SYSTEM]JUNK.SYS;
```

15.16 Changing Page, Swap, and Dump File Sizes

The following sections explain how to change sizes of page, swap, and dump files using different methods:

Method	For More Information
Using AUTOGEN (recommended method)	Section 15.16.1
Using SWAPFILES.COM (for primary files only)	Section 15.16.2
Using SYSGEN	Section 15.16.3

15.16.1 Using AUTOGEN (Recommended Method)

AUTOGEN automatically calculates appropriate sizes for page, swap, and dump files. It also modifies the files to the appropriate sizes and installs them. You can control sizes calculated by AUTOGEN by defining symbols in the file MODPARAMS.DAT. For more information, see Section 15.16.1.1.

How to Perform This Task

To change page, swap, and dump files, execute AUTOGEN in two passes as follows:

1. Enter the following command to invoke a first pass of AUTOGEN. In this pass, AUTOGEN displays its calculations for system file sizes to SYS\$OUTPUT:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS TESTFILES
```
2. If the file sizes displayed in step 1 are inadequate, add symbols to MODPARAMS.DAT to control the size of files as explained in Section 15.16.1.1 and return to step 1.
3. When you are satisfied with the file sizes displayed in step 1, execute a second pass of AUTOGEN using the following command to install the modified system files when the system is rebooted:

```
$ @SYS$UPDATE:AUTOGEN GENPARAMS REBOOT
```

15.16.1.1 Controlling the Size of Page, Swap, and Dump Files in MODPARAMS.DAT

You can add information to the AUTOGEN parameter file MODPARAMS.DAT to control the sizes that AUTOGEN calculates for page, swap, and dump files. If you do not supply system file size information in MODPARAMS.DAT, AUTOGEN performs default size calculations for page, swap, and dump files.

For information on AUTOGEN, see Section 14.4. For more information on MODPARAMS.DAT, see Section 14.4.4.

Managing System Page, Swap, and Dump Files

15.16 Changing Page, Swap, and Dump File Sizes

You can define symbols in MODPARAMS.DAT to specify either of the following:

Size to Be Specified	For More Information
Total desired size for all page or swap files on a system (not valid for the dump file)	Table 15-4
Sizes for <i>individual</i> page, swap, or dump files	Table 15-5

Note

You cannot specify sizes for both total and individual files. AUTOGEN issues a warning if conflicting symbol definitions exist in MODPARAMS.DAT.

For page and swap files, AUTOGEN generally manipulates the primary files SYS\$SYSTEM:PAGEFILE.SYS and SYS\$SYSTEM:SWAPFILE.SYS *only* if you have no other page and swap files; if you have secondary files, AUTOGEN manipulates the secondary files and excludes primary files. However, in some instances, AUTOGEN might modify the size of the primary page and swap files. If you do not want AUTOGEN to change the sizes of the primary files, specify the following symbols in MODPARAMS.DAT:

```
PAGEFILE = 0
SWAPFILE = 0
```

These symbols direct AUTOGEN to ignore the primary page and swap files when calculating sizes.

If the creation or extension of a file would cause the target disk to become more than 95 percent full, AUTOGEN issues a warning and does not perform the operation.

You can use AUTOGEN to create a page, swap, or dump file that is smaller than the current version of the file. After you have booted and begun using the new file, remember to use the DCL command PURGE to reclaim the disk space from the old version of the file. To determine the current sizes of installed page and swap files, enter the DCL command SHOW MEMORY/FILES. If you increased the size of any of these files and have not rebooted, this command displays the original sizes.

Note

AUTOGEN will not change file sizes if you specify a value of 0 or a value that is within 10 percent of the current size.

Table 15-4 lists the symbols you can define in MODPARAMS.DAT to control *total* size of page file, swap file, or dump file space.

Managing System Page, Swap, and Dump Files

15.16 Changing Page, Swap, and Dump File Sizes

Table 15–4 Symbols for Controlling the Total Size of Page, Swap, or Dump File Space

Operation	Page File Symbol	Swap File Symbol	Dump File Symbol
To define the total amount of space	PAGEFILE = n^1	SWAPFILE = n^1	DUMPFIL = n^1
To increase total size	ADD_PAGEFILE = n	ADD_SWAPFILE = n	ADD_DUMPFIL = n
To specify maximum total size	MAX_PAGEFILE = n	MAX_SWAPFILE = n	MAX_DUMPFIL = n
To specify minimum total size	MIN_PAGEFILE = n	MIN_SWAPFILE = n	MIN_DUMPFIL = n

¹ n is the total size, in blocks. If n is 0, the corresponding AUTOGEN section is skipped. For page and swap files, if n is not 0 and no secondary files exist, AUTOGEN applies the value to primary files. If n is not 0, and secondary files exist, AUTOGEN applies any change evenly across all secondary page or swap files but, in most cases, does not change primary files.

Table 15–5 lists the symbols you can define in MODPARAMS.DAT to control the size of *individual* files.

Table 15–5 Symbols for Controlling the Size of Individual Page and Swap Files

Operation	Page File Symbol ¹	Swap File Symbol ¹
To specify file size	PAGEFILE n _SIZE = <i>block-size</i>	SWAPFILE n _SIZE = <i>block-size</i>
To increase file size	ADD_PAGEFILE n _SIZE = <i>block-size</i>	ADD_SWAPFILE n _SIZE = <i>block-size</i>
To specify maximum file size	MAX_PAGEFILE n _SIZE = <i>block-size</i>	MAX_SWAPFILE n _SIZE = <i>block-size</i>
To specify minimum file size	MIN_PAGEFILE n _SIZE = <i>block-size</i>	MIN_SWAPFILE n _SIZE = <i>block-size</i>

¹For n , specify an integer that indicates the page or swap file. Refer to the primary page and swap files by specifying a value of 1 for n ; refer to subsequent files by specifying increasingly higher integer values for n . For example, to refer to a secondary page or swap file, specify a value of 2 for n . For *block-size*, specify the size in blocks.

Examples

The following line in MODPARAMS.DAT specifies that all page file space should total 100,000 blocks:

```
PAGEFILE = 100000
```

If you had only a primary page file, the resulting size of that file would be 100,000 blocks. If you had multiple page files, the difference between the total current size and the total new size would be spread across secondary files. For example, if you specified PAGEFILE = 100000, the changed page file sizes would be as follows:

File	Original Size (in Blocks)	Resulting Size (in Blocks)
Primary page file	10 000	10 000
Secondary page file 1	30 000	45 000
Secondary page file 2	30 000	45 000

To direct AUTOGEN to set the primary page file size to 10 000 blocks, use the symbol definition:

```
PAGEFILE1_SIZE = 10000
```


Managing System Page, Swap, and Dump Files

15.16 Changing Page, Swap, and Dump File Sizes

To direct AUTOGEN to create a new secondary swap file named PAGED\$:[PAGESWAP]SWAPFILE.SYS that holds 30 000 blocks, use the symbol definitions:

```
SWAPFILE2_NAME = "PAGED$:[PAGESWAP]SWAPFILE.SYS"
MIN_SWAPFILE2_SIZE = 30000
```

15.16.2 Using SWAPFILES.COM

Digital recommends you use AUTOGEN to change sizes of page, swap, and dump files. However, you can use the command procedure SYS\$UPDATE:SWAPFILES.COM to change the size of *primary* page, swap, and dump files. SWAPFILES.COM shows you the current size of the page, swap, and dump files before you change the sizes.

If you change the sizes of page, swap, or dump files, you must be sure to edit MODPARAMS.DAT to specify the new sizes, as explained in Section 15.16.1.1. If you do not specify the new sizes in MODPARAMS.DAT, AUTOGEN will resize the files next time it runs.

The procedure displays the sizes of the current page swap, and dump files in SYS\$SYSTEM, and the amount of space remaining on the system disk. It then allows you to enter new sizes, or keep the existing sizes for these files. If you specify a size that is larger than that of an existing file, the procedure automatically extends the size of a page or dump file. If you specify a smaller size for a system page, swap, or dump file, a new version of the file is created.

VAX

On VAX systems, if you are running from a saved Snapshot image, changing the size of any of the page, swap, or dump files disables the ability to boot from that image, and you must create a new Snapshot image. (For more information, see Section 4.7.) ♦

How to Perform This Task

1. Enter the following command to invoke the command procedure:

```
$ @SYS$UPDATE:SWAPFILES.COM
```

VAX

For VAX systems, if you are running from a saved Snapshot image, the following message is displayed

```
*****
* You are currently running from a saved snapshot image.          *
* If you change the size of any of the page, dump, or swapfiles, it *
* will be necessary to create a new snapshot image.              *
* Booting from the old snapshot image will be disabled.          *
*****
```

Do you wish to proceed [NO]?♦

The system displays the current files found in SYS\$SYSTEM and their sizes. For example:

Current file sizes are:

Directory SYS\$SYSROOT:[SYSEXE]

```
PAGEFILE.SYS;1    16384
SYSDUMP.DMP;1     4128
SWAPFILE.SYS;1    3072
```

Total of 3 files, 23584 blocks.

There are 128741 available blocks on SYS\$SYSDEVICE.

Managing System Page, Swap, and Dump Files

15.16 Changing Page, Swap, and Dump File Sizes

2. In response to the following prompt, type the desired size, in blocks, for the page file. To keep the same size, press Return:
Enter new size for page file:
3. In response to the following prompt, type the desired size, in blocks, for the dump file. To keep the same size, press Return:
Enter new size for system dump file:
4. In response to the following prompt, type the desired size, in blocks, for the swap file. To keep the same size, press Return:
Enter new size for swap file:
5. Shut down and reboot the system to use the new files.
6. After the system reboots, purge obsolete copies of the files. Do not delete the old files until the system reboots.
7. Edit MODPARAMS.DAT to include the new file sizes, as explained in Section 15.16.1.1. If you do not specify the new sizes in MODPARAMS.DAT, AUTOGEN will automatically resize the files the next time it runs.

Example

```
$ @SYS$UPDATE:SWAPFILES
```

To leave a file size at its current value type a carriage return in response to its size prompt.
Current file sizes are:

```
Directory SYS$SYSROOT:[SYSEXE]
```

```
PAGEFILE.SYS;1      100000
SYSDUMP.DMP;1       28000
SWAPFILE.SYS;1      33000
```

Total of 3 files, 161000 blocks.

There are 128741 available blocks on SYS\$SYSDEVICE.

Enter new size for page file:

Enter new size for system dump file: 30000

```
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]SYSDUMP.DMP;1 extended
```

Enter new size for swap file:

```
*****
* Please reboot in order for the new files to be used by the system. *
* After rebooting, purge obsolete copies of the files.                *
* DO NOT delete the old files until after the reboot.                  *
*****
```

15.16.3 Using SYSGEN

Digital recommends you use AUTOGEN to create and change page, swap, and dump files. AUTOGEN invokes the System Generation utility (SYSGEN) to create or change the files. However, in an emergency, you can use the System Generation utility (SYSGEN) to directly change the size of page, swap and dump files. For example, if you see that page file space is becoming dangerously low, you might use SYSGEN to quickly add page file space to prevent the system from hanging.

Managing System Page, Swap, and Dump Files

15.16 Changing Page, Swap, and Dump File Sizes

Note

The SWPFILCNT and PAGFILCNT system parameters limit the number of swap and page files that the system installs. See the *OpenVMS System Management Utilities Reference Manual* for more information.

How to Perform This Task

1. Determine the appropriate size of the files. For information, see Section 15.4.
2. Invoke SYSGEN and enter the CREATE command in the following format:

CREATE file-spec/SIZE=block-size

For *file-spec*, specify the full file specification.

For *block-size*, specify the size of the file in blocks.

If the file you specify already exists and the size you specify is larger than the existing file, the command extends the existing file. If the file you specify already exists and the size you specify is smaller than the existing file, the command creates a new file of the specified size.

For example, the following command extends the existing, smaller primary page file PAGEFILE.SYS:

```
SYSGEN> CREATE PAGEFILE.SYS/SIZE=100000
```

For more information on the SYSGEN command CREATE, see the SYSGEN section in the *OpenVMS System Management Utilities Reference Manual*.

Note

Frequent file creation and deletion can cause the free space on a disk to become severely fragmented. SYSGEN issues a HEADERFULL warning message if it determines that the creation or extension of a system file would cause that file to become fragmented enough to render the system unbootable. If this occurs, Digital recommends that you back up and restore your system disk to consolidate the free space on the volume into one contiguous area. (For more information, see Section 10.17.) After you have restored the disk, retry the SYSGEN operation. In cases where SYSGEN issues a warning message, the file might be somewhat larger, but not as large as the value specified in the CREATE command.

Managing System Page, Swap, and Dump Files

15.16 Changing Page, Swap, and Dump File Sizes

3. Use the following table to determine if you should reboot to use the new or modified file:

Type	Change	Reboot Required?
Primary page, swap, or dump file ¹	New file	Yes
	Extended file	Yes
Secondary page or swap file	New file	No ²
	Extended file	Yes

¹Primary page, swap, and dump files are SYS\$SPECIFIC:[SYSEXE] PAGEFILE.SYS, SWAPFILE.SYS, SYSDUMP.DMP.

²Although rebooting the system is unnecessary, you must install secondary files before the system can use them. For more information, see Section 15.14.

4. If you create a new version of the file, purge the old version *after* the system reboots.

Example

The commands in the following example extend the existing files PAGEFILE.SYS, SWAPFILE.SYS, and SYSDUMP.DMP to the specified sizes:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> CREATE PAGEFILE.SYS/SIZE=100000
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]PAGEFILE.SYS;1 extended
SYSGEN> CREATE SWAPFILE.SYS/SIZE=30000
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]SWAPFILE.SYS;1 extended
SYSGEN> CREATE SYSDUMP.DMP/SIZE=33000
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]SYSDUMP.DMP;1 extended
SYSGEN> EXIT
```


Performance Considerations

This chapter introduces the basic concepts of performance management. For more detailed information, see one of the following manuals:

- On VAX systems, see the *Guide to OpenVMS Performance Management*.
- On Alpha systems, see *A Comparison of System Management on OpenVMS AXP and OpenVMS VAX*.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Knowing your work load	Section 16.2
Choosing a workload management strategy	Section 16.3
Distributing the work load	Section 16.4
Predicting when tuning is required	Section 16.6
Evaluating tuning success	Section 16.7
Choosing performance options	Section 16.8
Installing images with the Install utility (INSTALL)	Section 16.9

This chapter explains the following concepts:

Concept	Section
Performance management	Section 16.1
System tuning	Section 16.5
Images and known images	Section 16.9.1
Known file lists	Section 16.9.2
Attributes of known images	Section 16.9.3

16.1 Understanding Performance Management

Performance management means optimizing your hardware and software resources for the current work load. This task entails several distinct but related activities:

- Acquiring a thorough familiarity with your work load and an understanding of how that work load exercises the system's resources. This knowledge, combined with an appreciation of the operating system's resource

Performance Considerations

16.1 Understanding Performance Management

management mechanisms, will enable you to establish realistic standards for system performance in areas such as the following:

- Interactive and batch throughput
- Interactive response time
- Batch job turnaround time
- Routinely monitoring system behavior to determine if, when, and why a given resource is approaching capacity.
- Investigating reports of degraded performance from users.
- Planning for changes in the system work load or hardware configuration and being prepared to make any necessary adjustments to system values.
- Performing, after installation, certain optional system management operations.

16.2 Knowing Your Work Load

One of the most important assets that a system manager brings to any performance evaluation is an understanding of the normal work load and behavior of the system. Each system manager must assume the responsibility for understanding the system's work load sufficiently to be able to recognize normal and abnormal behavior; to predict the effects of changes in applications, operations, or usage; and to recognize typical throughput rates. The system manager should be able to answer such questions as the following:

- What is the typical number of users on the system at any given time of day?
- What is the typical response time for various tasks for this number of users, at any given hour of operation?
- What are the peak hours of operation?
- Which jobs typically run at which time of day?
- Which commonly run jobs are intensive consumers of the CPU, memory, and disk space?
- Which applications involve the most image activations?
- Which parts of the system software, if any, have been modified or user-written, such as device drivers?
- Do any known, or anticipated, system bottlenecks exist?

If you are new to the OpenVMS operating system or to system management, you should observe system operation using the following tools:

- Monitor utility
- Accounting utility
- SHOW commands (available through DCL)

The *Guide to OpenVMS Performance Management* provides detailed procedures for using the Monitor utility and, to a lesser extent, other operating system tools to observe and evaluate system performance.

Over time you will learn about metrics such as the typical page fault rate for your system, the typical CPU usage, the normal memory usage, and typical modes of operation. You will begin to see how certain activities affect system performance and how the number of users or the time of day affects some of the values.

As you continue to monitor your system, you will come to know what range of values is acceptable, and you will be better prepared to use these same tools, together with your knowledge, to detect unusual conditions. Routine evaluation of the system is critical for effective performance management. The best way to avoid problems is to anticipate them; you should not wait for problems to develop before you learn how the system performs.

You can learn more about your system's operation if you use the Monitor and Accounting utilities on a regular basis to capture and analyze certain key data items. By observing and collecting this data, you will also be able to see usage trends and predict when your system may reach its capacity.

You should also understand that system resources are used by system management tools. Be careful, therefore, in selecting the items you want to measure and the frequency with which you collect the data. If you use the tools excessively, the consumption of system resources to collect, store, and analyze the data can distort your picture of the system's work load and capacity. The best approach is to have a plan for collecting and analyzing the data.

16.3 Choosing a Workload Management Strategy

System performance is directly proportional to the efficiency of workload management. Each installation must develop its own strategy for workload management. Before adjusting any system values, make sure to resolve the following issues and that your workload management strategy is correct:

- ☐ Does the work load "peak" at a particular time of day, that is, is it noticeably heavier than at other times?
- ☐ Can the work load be better balanced? Perhaps some voluntary measures can be adopted by users, after appropriate discussion.
- ☐ Could some jobs be run better as batch jobs, preferably during nonpeak hours?
- ☐ Have primary and secondary hours of operation been employed with users? If not, could system performance benefit by adopting this practice? If the primary and secondary hours are in use, are the choices of hours the most appropriate for all users? (Plan to review this issue every time you either add or remove users or applications, to ensure that the desired balance is maintained.)
- ☐ Can future applications be designed to work around any known or expected system bottlenecks? Can present applications be redesigned somewhat, for the same purpose? (See the *Guide to OpenVMS File Applications*.)
- ☐ Are you making the best use of the code-sharing ability that the operating system offers? Code sharing provides an excellent means to conserve memory and improve performance over the life of the system.

16.4 Distributing the Work Load

You should distribute the work load as evenly as possible over the time your system is running. Although the work schedule for your site may make it difficult to schedule interactive users at optimum times, the following techniques may be helpful:

- Run large jobs as batch jobs—Establish a site policy that encourages the submission of large jobs on a batch basis. Regulate the number of batch streams so that batch usage is high when interactive usage is low. You might also want to use DCL command qualifiers to run batch jobs at lower priority, adjust the working set sizes, or control the number of concurrent jobs. For information about setting up your batch environment, see Section 13.5.
- Restrict system use—Do not permit more users to log in at one time than the system can support with an adequate response time. You can restrict the number of interactive users with the DCL command SET LOGINS /INTERACTIVE. You can also control the number of concurrent processes with the MAXPROCESSCNT system parameter, and the number of remote terminals allowed to access the system at one time with the RJOBLIM system parameter. See Section 14.5 for information about modifying system parameters. See the *OpenVMS System Management Utilities Reference Manual* for descriptions of all system parameters.

You might also restrict use of the system by groups of users to certain days and hours of the day. You can use the Authorize utility to define the permitted login hours for each user. In particular, refer to the AUTHORIZE qualifiers /PRIMEDAYS, /P_RESTRICT, /PFLAGS, /SFLAGS, and /S_RESTRICT. For more information, see Chapter 6 and the AUTHORIZE section of the *OpenVMS System Management Utilities Reference Manual*.

You can use the DCL command SET DAY to override the conventional day of the week associations for primary and secondary days. For example, you might want to specify a primary day of the week as a secondary day when it is a holiday.

- Design applications to reduce demand on binding resources—If you know where your system bottlenecks are or where they will likely occur in the near future, you can distribute the work load more evenly by planning usage that minimizes demand on any bottleneck points. (See the *Guide to OpenVMS File Applications*.)

16.5 Understanding System Tuning

Tuning is the process of altering various system values to obtain the optimum overall performance possible from any given configuration and work load. However, the process does not include the acquisition and installation of additional memory or devices, although in many cases such additions (when made at the appropriate time) can vastly improve system operation and performance.

Always aim for best overall performance, that is, performance viewed over time. The work load is constantly changing on most systems. System parameters that produce optimal performance at one time may not produce optimal performance a short time later as the work load changes. Your goal is to establish values that, on average, produce the best overall performance.

Performance Considerations

16.5 Understanding System Tuning

Before you undertake any action, you must recognize that the following sources of performance problems cannot be cured by adjusting system values:

- Improper operation
- Unreasonable performance expectations
- Insufficient memory for the applications attempted
- Inadequate hardware configuration for the work load, such as too slow a processor, too few buses for the devices, too few disks, and so forth
- Improper device choices for the work load, such as using disks with insufficient speed or capacity
- Hardware malfunctions
- Human errors, such as poor application design or allowing one process to consume all available resources

When you make adjustments, you normally select a very small number of values for change, based on a careful analysis of the behavior being observed. You control system resources by tuning the values of two types of parameters:

Parameter Type	Description
System parameters	<p>The values set for system parameters control system resources on a systemwide basis. The AUTOGEN command procedure automatically sets system parameters to appropriate values for your system configuration. AUTOGEN can also record feedback from a running system to adjust those parameters based on the system's work load. The <i>Guide to OpenVMS Performance Management</i> describes how to select the parameters and new values that are likely to produce the desired changes.</p> <p>Section 14.5 explains how to use AUTOGEN to modify system parameter values.</p>
UAF limits and quotas	<p>The values set for limits and quotas in each User Authorization File (UAF) record control system resources on a per-user basis. To control these values, use the Authorize utility. For information, see Section 6.11.</p>

Before you undertake any tuning operation, be sure you are familiar with the resource management mechanisms described in the *Guide to OpenVMS Performance Management* and *A Comparison of System Management on OpenVMS AXP and OpenVMS VAX*. Understand the nature of system values before adjusting them. Without the proper level of understanding, you might very well degrade, rather than improve, overall performance.

Finally, while investigating the cause of an apparent performance problem, keep in mind that tuning is a last resort.

16.6 Predicting When Tuning Is Required

Under most conditions, tuning is rarely required for OpenVMS systems. The AUTOGEN command procedure, which is included in the operating system, establishes initial values for all the configuration-dependent system parameters so that they match your particular configuration. For information about AUTOGEN, see Section 14.4.

Performance Considerations

16.6 Predicting When Tuning Is Required

Additionally, the system includes features that, in a limited way, permit it to adjust itself dynamically during operation. That is, the system detects the need for adjustment in certain areas, such as the nonpaged dynamic pool, the working set size, and the number of pages on the free and modified page lists. The system makes rough adjustments in these areas automatically. As a result, these areas can grow dynamically, as appropriate, during normal operation.

Experience has shown that the most common cause of disappointment in system performance is insufficient hardware capacity. Once the demand on a system exceeds its capacity, adjusting system values will not result in any significant improvements, simply because such adjustments are a means of trading off or juggling existing resources.

Although tuning is rarely required, you should recognize that system tuning may be needed under the following conditions:

- If you have adjusted your system for optimal performance with current resources and then acquire new capacity, you must plan to compensate for the new configuration. In this situation, the first and most important action is to execute the AUTOGEN command procedure.

For more information about AUTOGEN, see Section 14.4.

- If you anticipate a dramatic change in your work load, you should expect to compensate for the new work load.

16.7 Evaluating Tuning Success

Whenever you adjust your system, you should monitor its behavior afterward to be sure that you have obtained the desired results. To observe results, use the Monitor utility and the various forms of the DCL command SHOW. See the *OpenVMS DCL Dictionary* for detailed information on the SHOW command. See Section 18.8.2 for information about using MONITOR. See the *OpenVMS System Management Utilities Reference Manual* (MONITOR) for detailed descriptions of MONITOR commands.

For example, you might consider running some programs whose results you believe are fixed and reproducible at the same time that you run your normal work load. If you run the programs and measure their running times under nearly identical workload conditions both before and after your adjustments, you can obtain a basis for comparison.

However, when applying this technique, remember to take the measurements under very similar workload conditions. Also, remember that this test alone does not provide conclusive proof of success. The possibility always exists that your adjustments may have favored the performance of the image you are measuring—to the detriment of other images. Therefore, in all cases, continue to observe system behavior closely for a time after you make any changes.

16.8 Choosing Performance Options

Following is a list of optional system management operations, normally performed after installation, that often result in improved overall performance. Choose the options that are appropriate for your site. Not all options are appropriate at every site.

- ☐ Decompress system libraries—Most of the libraries shipped with the operating system are in a compressed format in order to conserve disk space. The system dynamically decompresses them whenever they are

accessed, and the resulting performance slowdown is especially noticeable during link operations and when requesting online help. If you have sufficient disk space, decompressing the libraries improves both CPU and elapsed time performance. To do this, invoke the command procedure `SYS$UPDATE:LIBDECOMP.COM`. The decompressed object libraries use about 25 percent more disk space than when compressed; the decompressed help libraries use about 50 percent more disk space.

- ☐ Disable file system high-water marking—This security feature is set by default when a volume is initialized to guarantee that users cannot read data they have not written.

For non-shared sequential files, the performance impact of high-water marking is minimal. However, for files of non-sequential format, high-water marking creates some overhead; the system erases the previous contents of the disk blocks allocated every time a file is created or extended.

Disabling the feature improves system performance by a variable amount, depending on the following factors:

- How frequently new files are created
- For indexed and relative files, how frequently existing files are extended
- How fragmented the volume is

Be sure to consider the security implications before you disable high-water marking.

To disable high-water marking, you can specify the `/NOHIGHWATER` qualifier when initializing the volume, or you can disable high-water marking with the DCL command `SET VOLUME` in the following format:

`SET VOLUME/NOHIGHWATER_MARKING device-spec[:]`

- ☐ Set file extend parameters for OpenVMS Record Management Services (RMS)—Because files extend in increments of twice the multiblock count (default 16), system defaults provide file extension of 32 blocks rounded up to the nearest multiple of the disk's cluster size. Thus, when files are created or extended, increased I/O may slow performance. The problem can be corrected by specifying larger values for file extend parameters or by setting the system parameter `RMS_EXTEND_SIZE`. See Section 14.5 for information about modifying system parameters. See the *OpenVMS System Management Utilities Reference Manual* for a description of all system parameters.

For more information about establishing the file extension quantity, see the section on tuning in the *Guide to Creating OpenVMS Modular Procedures*.

VAX

- ☐ On VAX systems, relink images—Beginning with VAX/VMS Version 4.0, the Run-Time Library (VMSRTL) was separated into five smaller libraries. Running images linked under previous versions of the operating system will therefore incur the image activation costs of mapping all five libraries, even if only one is needed. You may improve performance by relinking pre-Version 4.0 images that reference run-time library routines, so that only the required libraries are mapped and activated. ♦
- ☐ Install frequently used images—When an image is accessed concurrently by more than one process on a routine basis, install the image with the `INSTALL` utility, specifying the `/OPEN`, `/SHARED`, and `/HEADER_RESIDENT` qualifiers. You will thereby ensure that all processes use the same physical copy of the image, and that the image will be activated in the most efficient way.

Performance Considerations

16.8 Choosing Performance Options

Generally, an image takes about two additional physical pages when installed with the /OPEN, /HEADER_RESIDENT, and /SHARED qualifiers. The utility's LIST/FULL command shows the highest number of concurrent accesses to an image installed with the /SHARED qualifier. This information can help you decide whether installing an image is an efficient use of memory.

See Section 16.9.11 and the INSTALL section of *OpenVMS System Management Utilities Reference Manual* for more information on installing images.

Alpha

- ☐ On Alpha systems, install shareable and executable images specifying the /RESIDENT qualifier with the Install utility. For more information, see Section 16.9.6.

Note that this is a tradeoff between the CPU and memory. Installing an image with /RESIDENT qualifier means that the code is to be nonpaged. Depending on the amount of sharing, this can be a memory gain or loss. ♦

- ☐ Reduce system disk I/O—You can move frequently accessed files off the system disk and use logical names to specify the location or, where necessary, other pointers to access them. For example:
 - SYSUAF.DAT (SYSUAF is the logical name)
 - RIGHTSLIST.DAT (RIGHTSLIST is the logical name)
 - VMSMAIL_PROFILE.DATA (VMSMAIL is the logical name)
 - NETPROXY.DAT (NETPROXY is the logical name)
 - NET\$PROXY.DAT (NET\$PROXY is the logical name)
 - The queue database (for more information, see Section 12.3)
 - ERRFMT log files (SYS\$ERRORLOG is the logical name)
 - MONITOR log files (SYS\$MONITOR is the logical name)
 - The accounting log file (ACCOUNTNG is the logical name)
 - SECURITY_AUDIT.AUDIT\$JOURNAL (SET AUDIT /JOURNAL=SECURITY/DESTINATION= *filespec*)
 - Default DECnet for OpenVMS accounts (records included in the SYSUAF file on the OpenVMS distribution kit)

To redefine logical names for these system files, edit the site-specific command procedure SYS\$MANAGER:SYLOGICALS.COM. For more information on defining logical names in SYLOGICALS.COM, see Section 5.2.5.

You can also consider moving paging and swapping activity off the system disk by creating large secondary page and swap files on a less heavily used disk. However, if you want to store crash dumps for diagnosing system failures, the dump file must reside in the system-specific directory SYS\$SPECIFIC:[SYSEXEC] on the system disk for storing crash dumps; if no dump file exists in SYS\$SPECIFIC:[SYSEXEC], the primary page file must be located there if you want to store crash dumps. For detailed information on moving page and swap files, see Section 15.13.

16.9 Using INSTALL to Install Known Images

The Install utility (INSTALL) stores information about images in memory. Use INSTALL for the following reasons:

Reason	For More Information
To conserve memory use for images that are used concurrently	Section 16.9.4
To improve system performance	Section 16.9.5
‡On Alpha systems, with sliced images to improve performance	Section 16.9.6
To make programs that require enhanced privileges available for general use	Section 16.9.7
To allow a nonprivileged process to perform the privileged functions of the image	Section 16.9.7
To mark a sharable image as trusted so it can be invoked by privileged executable images	Section 16.9.7
‡Alpha specific	

The site-independent startup command procedure, STARTUP.COM, uses INSTALL to install certain system images when the system boots. You use INSTALL to install other selected images, according to the needs of your site.

Installed images must be reinstalled each time the system reboots. To do so, include INSTALL commands in the site-specific startup command procedure SYSTARTUP_VMS.COM, as explained in Section 5.2.7.

Note that Install utility (INSTALL) commands perform a different function than System Generation utility (SYSGEN) INSTALL commands.

The following sections explain installed images and how to use the Install utility.

16.9.1 Understanding Images and Known Images

An **image** is a collection of procedures and data bound together by the Linker utility to form an executable program. Executable programs can be executed (or run) by a process. Usually, executable programs have the file type .EXE.

There are three types of images:

Image Type	Description
Executable	An image linked with the /EXECUTABLE qualifier (or without the /SHAREABLE qualifier) of the Linker utility. For more information, see the <i>OpenVMS Linker Utility Manual</i> .
Shareable	An image linked with the /SHAREABLE qualifier of the Linker utility; it must subsequently be linked into an executable image to be used. (Shareable images are sometimes referred to as linkable images, because they can be specified—implicitly or explicitly—as input files to the link of another file.) A shareable image is not copied into the executable images that link with it. Thus, only one copy of the shareable image needs to be on disk, no matter how many executable images have linked with it. For more information, see the <i>OpenVMS Linker Utility Manual</i> .

Performance Considerations

16.9 Using INSTALL to Install Known Images

Image Type	Description
System	An image that does not run under the control of the operating system. It is intended for standalone operation only. The content and format of a system image differs from that of shareable images and executable images. For more information, see the <i>OpenVMS Linker Utility Manual</i> .

When you install an image with INSTALL, the image is assigned attributes and becomes known to the system. For this reason, an installed image is also called a **known image**.

The DCL command RUN parses search lists in a manner that favors known images. On its first pass through the search list, the RUN command looks up images on known file lists and executes each known image that it finds. On its second pass through the search list, the RUN command looks up images on disk and executes those images not executed in the first pass.

16.9.2 Understanding Known File Lists

The system defines known images in internal data structures called **known file lists**. Each entry in the known file list identifies the file name of the installed file and the attributes with which it was installed (for information about attributes of installed images, see Section 16.9.3).

A separate known file list exists for all installed images whose device, directory, and file type are identical. For example, all installed images with the file name DISK\$VOLUME:[MAIN]filename.EXE would be in one known file list, and all installed images with the file name DISK\$VOLUME:[TEST]filename.EXE would be in another known file list.

Known file lists last only while the system is operating. If the system is shut down or fails for any reason, you must reinstall all known images after the system is rebooted.

16.9.3 Understanding Attributes You Can Assign to Known Images

By specifying appropriate qualifiers to INSTALL commands, you can assign attributes to known images. Table 16-1 describes these attributes and the qualifiers that are used to assign them to known images.

Table 16-1 Attributes of Known Images

Attribute	Description	Qualifier
Header resident	The header of the image file (native images only) remains permanently resident, saving one disk I/O operation per file access. For images with single-block file headers, the cost is less than 512 bytes of paged dynamic memory per file; for images with multiblock headers, the cost varies according to the header block count. The images must also be declared permanently open.	/[NO]HEADER_RESIDENT
Permanently open	Directory information on the image file remains permanently resident, eliminating the usual directory search required to locate a file. The cost of keeping an image file permanently open is approximately 512 bytes of paged dynamic memory per file.	/OPEN

(continued on next page)

Table 16-1 (Cont.) Attributes of Known Images

Attribute	Description	Qualifier
Privileged	Amplified privileges are temporarily assigned to any process running the image, permitting the process to exceed its user authorization file (UAF) privilege restrictions during execution of the image. In this way, users with normal privileges can run programs that require higher-than-normal privileges. This attribute (and the /PRIVILEGED qualifier that creates it) applies only to executable images.	/PRIVILEGED[=(privilege,...)]
Protected	A shareable image contains protected code, that is, code that runs in kernel or executive mode but that can be called by a user-level image. Protected images must be declared shareable.	/PROTECTED
‡Resident	On Alpha systems, improves the performance of shareable or executable images that have been linked with /SHARE and a new LINK qualifier, /SECTION_BINDING=(CODE,DATA), by installing them as resident with the Install utility. The code sections of an installed resident shareable image reside in huge pages called granularity hint regions (GHRs) in memory. The Alpha hardware can consider a set of pages as a single GHR. This GHR can be mapped by a single page table entry (PTE) in the translation buffer (TB). The result is a reduction in TB miss rates. For more information, see Section 16.9.6.	/RESIDENT
Shareable	More than one user can access the read-only and non-copy-on-reference read/write sections of the image concurrently, so that only one copy of those sections ever needs to be in physical memory. (Copy-on-reference sections always require a separate copy for each process.) The image is implicitly declared permanently open.	/SHARED
Writable	When a shareable non-copy-on-reference writable section is removed from physical memory (for paging reasons or because no processes are referencing it), it is written back to the image file. Any updates made by processes mapped to the section, therefore, are preserved (while the initial values are lost). The image must also be declared shareable.	/WRITABLE
‡Alpha specific		

16.9.4 Installing Images to Conserve Memory

Shareable images conserve memory because only one copy of the code needs to be in memory at any time, and many users can access the code concurrently. The Install utility is the only way to install images. Use the /SHARED qualifier to install images as shareable images.

When you install a shareable image, more than one user can access the read-only and non-copy-on-reference read/write sections of the image concurrently, so that only one copy of those sections ever needs to be in physical memory. (Copy-on-reference sections always require a separate copy for each process.) The image is implicitly declared permanently open.

When you install an image with the shareable attribute, permanent system global sections are created. Execution of non-copy-on-reference global sections requires only one copy per section to be in physical memory, no matter how many processes are running the image to which the sections belong.

Performance Considerations

16.9 Using INSTALL to Install Known Images

The number of images you can install with the shareable attribute is restricted by the GBLPAGES and GBLSECTIONS system parameters. For more information on these system parameters, see the *OpenVMS System Management Utilities Reference Manual*.

When an image is not installed, or is installed without the shareable attribute, each process running the image requires private sections in memory.

A shareable image linked to an executable image does not have to be installed to be executed. At image execution time, the system creates private sections from the shareable image. The only exception is that a shareable image containing a writable non-copy-on-reference section must be installed as a known image with the shareable and writable attributes.

16.9.5 Installing Images to Improve Image Performance

Image performance improves when programs are installed because the operating system opens any installed file by file ID rather than by file name, thus eliminating costly directory operations.

Installing images as header resident further enhances performance because the system avoids the overhead of I/O operations to read the image header into memory.

Note

On VAX systems, virtual I/O cache can automatically improve image performance at a level similar to that gained by installing images. However, you should decide whether to install it based on the configuration and requirements of your site. For more information on virtual I/O cache, see the *Guide to OpenVMS Performance Management*.

To install an image as header resident, specify the `/HEADER_RESIDENT` qualifier when you install the image. This makes the header of the image file (native images only) remain permanently resident, saving one disk I/O operation per file access. For images with single-block file headers, the cost is less than 512 bytes of paged dynamic memory per file; for images with multiblock headers, the cost varies according to the header block count. The images must also be declared permanently open by specifying the `/OPEN` qualifier.

Frequently accessed images, critical to a site's operations, can be installed as open images. To install an image as permanently open, specify the `/OPEN` qualifier when you install the image. This makes the directory information on the image file remain permanently resident, eliminating the usual directory search required to locate a file. The cost of keeping an image file permanently open is approximately 512 bytes of paged dynamic memory per file.

16.9.6 Installing Resident Images to Improve Performance (Alpha Only)

Alpha

On Alpha systems, you can improve the performance of shareable images that have been linked with `/SHARE` and a new `LINK` qualifier, `/SECTION_BINDING=(CODE,DATA)`, by installing them as resident with the `Install` utility. The code sections of an installed resident shareable image reside in huge pages called granularity hint regions (GHRs) in memory. The Alpha hardware can consider a set of pages as a single GHR. This GHR can be mapped by a single page table entry (PTE) in the translation buffer (TB). The result is a reduction in TB miss rates.

Performance Considerations

16.9 Using INSTALL to Install Known Images

This feature enables the operating system to split the contents of images and sort the pieces so that they can be placed with other pieces that have the same page protection in the same area of memory. Consequently, TBs on Alpha systems are used more efficiently than if the images were loaded in the traditional manner.

Application programmers are the likely users of the slicing feature for shareable images. As system manager, you might be asked to coordinate or assist slicing efforts by installing images as resident shareable images. Specify the `/RESIDENT=(CODE,DATA)` qualifier with the `INSTALL` commands `ADD`, `CREATE`, and `REPLACE` to install shareable and executable images as resident.

Resident images can also be installed with shareable linkage sections. The user has no direct control over which images are installed with shareable linkage sections. Images that are eligible for sharing linkage sections are:

- `CMA$TIS_SHR.EXE`
- `DECC$SHR.EXE`
- `DPML$SHR.EXE`
- `LIBOTS.EXE`
- `LIBRTL.EXE`

Linkage data for these images will be shared if space is provided for images in the process control region. Allocation of this space is governed by the `IMGREG_PAGES` system parameter. By default, adequate space is provided for the five images. Shared linkage reduces image activation time and decreases demand for physical memory.

You cannot remove images installed with shareable linkage from the known image list nor can you replace them, except by rebooting the system. To disable shareable linkage sections, set the system parameter `IMGREG_PAGES` to 0. ♦

16.9.7 Installing Images to Enhance Privileges of Images

There are two ways to allow an image to execute in an enhanced privilege environment:

- Installing existing executable images with extra privileges to allow a nonprivileged process to perform the privileged functions of the image. Use the `/PRIVILEGED` qualifier for the `INSTALL` commands `ADD` or `CREATE`.
- Installing privileged shareable images (which are used to implement user-written system services), allowing other, nonprivileged images to execute select portions of privileged code without enhancing the privileges of those individual images. Use the `/PROTECTED` and `/SHARED` qualifiers for the `INSTALL` commands `ADD` or `CREATE`.

Caution

Installing an image with enhanced privilege can compromise system security. Make sure the image does not enable a user to regain control with extra privileges enabled.

Performance Considerations

16.9 Using INSTALL to Install Known Images

16.9.7.1 Privileged Executable Images

A nonprivileged process can perform the privileged functions of an executable image when it is installed as a privileged image. Install executable images with enhanced privileges by using the `/PRIVILEGED` qualifier; amplified privileges are temporarily assigned to any process running the image (executable images only), permitting the process to exceed its user authorization file (UAF) privilege restrictions during execution of the image. In this way, users with normal privileges can run programs that require higher-than-normal privileges.

For an image installed with privileges to activate another image, such as a shareable image, either by having it linked to the privileged image or by using `LIB$FIND_IMAGE_SYMBOL`, the following conditions hold:

- The shareable image must be installed as a known image using `INSTALL`.

Note

Installing the Install utility itself requires that a number of shareable images have been previously installed. If any of those required shareable images (such as `SMG$SHR`, `LIBOTS`, and so on) is unavailable, the execution of the Install utility fails. Since `INSTALL` will not work in this situation, you cannot simply install the missing images. To work around this problem, redefine the `INSTALL` command as follows:

```
$ DEFINE INSTALL SYS$SYSTEM:INSTALL.EXE;0
```

When you now enter the `INSTALL` command, the image activator does not check the known files list for `INSTALL.EXE`, and the `INSTALL` command will complete, allowing you to install the required shareable images.

-
- Logical names and table names used to find the image must be defined in executive or kernel mode. In particular, the standard executive-mode definition of `LN$FILE_DEV` translates only to `LN$SYSTEM`; definitions in the process, job, or group tables are not recognized.
 - Only images linked with the Linker utility qualifiers `/NODEBUG` and `/NOTRACE` can be installed with enhanced privilege.

16.9.7.2 Privileged Shareable Images

A user-written system service assumes the privileges it requires when you install it as a privileged shareable image. To create a privileged, shareable image, you must:

1. Link a shareable image with the `/PROTECT` command qualifier or the `PROTECT=` option of the Linker utility, so that the image acquires its particular form of enhanced privileges.
 - Use the `/PROTECT` command qualifier when all parts of an image require protection.
 - Use the `PROTECT=` option when only part of a privileged shareable image requires protection.
2. Install the privileged shareable image with the Install utility, specifying both the `/PROTECTED` and the `/SHARED` qualifiers. The `/PROTECTED` qualifier assigns the protected attribute. The `/SHARED` qualifier assigns the shareable attribute. See Section 16.9.3 for information about these attributes.

Performance Considerations

16.9 Using INSTALL to Install Known Images

Note

You cannot create a privileged shareable image using the /PRIVILEGED qualifier for the INSTALL commands ADD or CREATE. This qualifier works only for executable images.

For more information on creating privileged shareable images, see the *OpenVMS Programming Concepts Manual*.

16.9.8 Installing Images to Allow Execution of Images Without Read Access

When a process runs an executable or shareable image to which it has execute but not read access, the image activator enters a restricted mode of operation similar to that entered when a privileged program is run. In this mode of operation:

- All shareable images activated during the life of the execute-only image must be installed.
- The image activator directs OpenVMS RMS to use only **trusted logical names** (logical names associated with executive or kernel mode) when opening image files.

Note

The executable image that calls an execute-only shareable image must be installed with the /EXECUTE_ONLY qualifier, which enables the executable image to activate shareable images to which the process has execute but not read access.

The /EXECUTE_ONLY qualifier has meaning only for executable images.

16.9.9 Determining Which Images to Install

You should install images that meet the following conditions:

- Images that run frequently
- Images that usually run concurrently from several processes
- Images that require special privileges
- On Alpha systems, images that have been linked with the Linker utility qualifier /SECTION_BINDING=(CODE,DATA)

Alpha

You can use ANALYZE/IMAGE on an Alpha system to determine whether an image is linked with /SECTION_BINDING=(CODE,DATA). In the ANALYZE/IMAGE output, look for the EIHD\$V_BIND_CODE or the EIHD\$V_BIND_DATA symbol; a value of 1 indicates that the /SECTION_BINDING=CODE or the /SECTION_BINDING=DATA qualifier was used, respectively. For more information, see the *OpenVMS Linker Utility Manual*. ♦

Because an installed file requires system resources, such as paged dynamic memory, install those files that most improve system performance and site requirements. The INSTALL command LIST provides information about installed images to help you evaluate the merits of installing images. For example, the LIST command calculates the number of times each image is accessed, and shows the number of concurrent accesses, so you can determine if the installation of the images is worth the overhead.

Performance Considerations

16.9 Using INSTALL to Install Known Images

16.9.10 Specifying File Names in INSTALL

When you use INSTALL commands, your file specifications must name existing executable or shareable images. OpenVMS Record Management Services (RMS) resolves each file specification using the following defaults:

- A device and directory type of SYS\$SYSTEM
- A file type of .EXE

Unless a file shares these defaults, you must specify a device and directory name and a file type with each file name. The highest existing version of the file is used by default. However, you can specify another version of the file as the known version of the image. Even if other versions of the file exist, the version that you specify will be the version that satisfies all known file lookups for the image.

16.9.11 Installing Images with INSTALL

Before performing this task, you should understand the following:

- Attributes of installed images. For information, see Section 16.9.3.
- File specifications for the Install utility. For information, see Section 16.9.10.

How to Perform This Task

1. Give yourself the CMKRNL privilege by entering the following command:

```
$ SET PROCESS/PRIVILEGES=CMKRNL
```

2. Invoke INSTALL by entering the following command:

```
$ INSTALL
```

3. Enter the ADD command in the following format:

```
ADD file-spec [/qualifier...]
```

Specify one or more of the following qualifiers, depending on which attributes you want to assign to the image:

```
/EXECUTE_ONLY  
/HEADER_RESIDENT  
/OPEN  
/PRIVILEGED  
/PROTECTED  
/RESIDENT (Alpha systems only)  
/SHARED  
/WRITABLE
```

For more information on installing images, see the INSTALL command ADD in the INSTALL section of the *OpenVMS System Management Utilities Reference Manual*.

16.9.12 Displaying Known Images with INSTALL

Use the INSTALL command LIST to display information about known images.

The information displayed with the /FULL qualifier of the LIST command can help you determine if installing an image is worth the expense.

Performance Considerations

16.9 Using INSTALL to Install Known Images

How to Perform This Task

1. Invoke INSTALL by entering the following command:

```
$ INSTALL
```

2. To display a list of all known images and their attributes, enter the LIST command as follows:

```
INSTALL> LIST
```

To display attributes for a specific image, specify the name of the image as follows:

```
LIST filename
```

For example:

```
INSTALL> LIST LOGINOUT
```

To display complete information about a specific image, including the number of accesses, the number of concurrent accesses, and the number of global sections created, specify the /FULL qualifier as follows:

```
LIST/FULL filename
```

Example

The following example displays complete information about the installed image LOGINOUT.EXE, including the number of accesses, the number of concurrent accesses, and the number of global sections created:

```
$ INSTALL
INSTALL> LIST/FULL LOGINOUT
DISK$VMS551:<SYS2.SYSCOMMON.SYSEXE>.EXE
LOGINOUT;2      Open Hdr Shar Prv
  Entry access count      = 36366
  Current / Maximum shared = 1 / 10
  Global section count    = 3
  Privileges = CMKRNLSYSNAMLOG_IOALTPRITMPMBXSYSPRV
INSTALL>
```

16.9.13 Defining Logical Names for Shareable Image Files

If a shareable image is not located in SYS\$SHARE, you must define a logical name for that image in order to run an executable image linked against it. For example, if the file specification for STATSHR is SYS\$SHARE:STATSHR.EXE, no logical name is necessary. But if you put STATSHR in SYS\$DEVICE:[TEST], you must define STATSHR as a logical name before running an executable image that calls it. The logical name must be the same one that was used as the input file specification for the shareable image when it was linked (this is the same name used in installation). For example:

```
$ DEFINE STATSHR SYS$SYSDEVICE:[TEST]STATSHR
```

By redefining the logical name of a shareable image, you can replace that shareable image with another without requiring the calling executable image to relink. For example, the following statement redefines the file name STATSHR. It becomes the logical name of the shareable image SYS\$SYSDEVICE:[MAIN]STATSHR.EXE for executable images calling STATSHR.

Performance Considerations

16.9 Using INSTALL to Install Known Images

```
$ DEFINE STATSHR SYSS$SYSDEVICE:[MAIN]STATSHR
```

Note

Logical names defined in the process or group logical name table are ignored when you run a privileged executable image. Only logical names and table names defined in executive or kernel modes are used to find the image.

Two shareable images installed with the /SHARED qualifier cannot have the same file name. (Use the INSTALL command REPLACE to update file versions.) For more information on the INSTALL command REPLACE, see the INSTALL section of the *OpenVMS System Management Utilities Reference Manual*.

16.9.14 Removing Known Images

The INSTALL command DELETE removes a known file list entry for an image and deletes all global sections created when the image was installed. Note the following restrictions on removing known images:

- A known image is not deleted as soon as the INSTALL DELETE command is entered. The deletion occurs only after all processes using the image have released it.
- A volume cannot be dismounted while any known file lists associated with it contain entries. To dismount a volume, you must delete all known images associated with it. You must also wait for all processes using those images to release them and for the system to write writable images back to their files. Use the DCL command SHOW DEVICES/FILES to determine the status of the files.

For more information on the INSTALL command DELETE, see the INSTALL section of the *OpenVMS System Management Utilities Reference Manual*.

Testing the System with UETP

This chapter explains how to use UETP (user environment test package) to test whether the OpenVMS operating system is installed correctly.

17.1 Overview

This overview summarizes what UETP does and how you use it. The rest of the chapter provides detailed instructions for setting up your system for testing, running the tests, and troubleshooting errors.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Running UETP (a summary)	Section 17.1.2
Preparing to use UETP	Section 17.2
Setting up the devices to be tested	Section 17.3
Starting UETP	Section 17.4
Stopping a UETP operation	Section 17.5
Troubleshooting: identifying and solving problems	Section 17.7

This chapter explains the following concepts:

Concept	Section
Understanding UETP	Section 17.1.1
Troubleshooting (an overview)	Section 17.6
UETP Tests and Phases	Section 17.8

17.1.1 Understanding UETP

UETP is a software package designed to test whether the OpenVMS operating system is installed correctly. UETP puts the system through a series of tests that simulate a typical user environment by making demands on the system that are similar to demands that can occur in everyday use.

UETP is not a diagnostic program; it does not attempt to test every feature exhaustively. When UETP runs to completion without encountering nonrecoverable errors, the system being tested is ready for use.

UETP exercises devices and functions that are common to all OpenVMS systems, with the exception of optional features such as high-level language compilers. The system components tested include the following:

- Most standard peripheral devices

Testing the System with UETP

17.1 Overview

- System's multiuser capability
- DECnet for OpenVMS software
- Clusterwide file access and locks

17.1.2 Summary of How to Use UETP

This section summarizes the procedure for running all phases of UETP with default values. If you are familiar with the test package, refer to this section. If you want additional information, refer to Section 17.2.

Alpha

Note

If you are using UETP on an OpenVMS Alpha system, you must execute the `CREATE_SPECIAL_ACCOUNTS.COM` command procedure to create the `SYSTEST` and `SYSTEST_CLIG` accounts before you begin the following procedure. For complete information about the `CREATE_SPECIAL_ACCOUNTS.COM` command procedure, see Section 6.4. ♦

1. Log in to the `SYSTEST` account as follows:

Username: `SYSTEST`
Password:

Caution

Because the `SYSTEST` and `SYSTEST_CLIG` accounts have privileges, unauthorized use of these accounts can compromise the security of your system.

2. Make sure no user programs are running and no user volumes are mounted.

Caution

By design, UETP assumes and requests the exclusive use of system resources. If you ignore this restriction, UETP can interfere with applications that depend on these resources.

3. After you log in, check all devices to be sure that the following conditions exist:
 - All devices you want to test are powered up and are on line to the system.
 - Scratch disks are mounted and initialized.
 - Disks contain a directory named `[SYSTEST]` with `OWNER_UIC=[1,7]`. (You can create this directory with the DCL command `CREATE /DIRECTORY`.)
 - Scratch magnetic tape reels are *physically* mounted on each drive you want tested and are initialized with the label `UETP` (using the DCL command `INITIALIZE`). Make sure magnetic tape reels contain at least 600 feet of tape.
 - Scratch tape cartridges have been inserted in each drive you want to test and are initialized with the label `UETP` (using the DCL command `INITIALIZE`).

- Line printers and hardcopy terminals have plenty of paper.
- Terminal characteristics and baud rate are set correctly. (See the user's guide for your terminal.)

Note that some communication devices need to be set up by Digital Services. (See Section 17.3.)

If you encounter any problems in preparing to run UETP, read Section 17.3 before proceeding.

4. To start UETP, enter the following command and press Return:

```
$ @UETP
```

UETP responds with the following question:

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?
```

Press Return to choose the default response enclosed in brackets. UETP responds with the following sequence of questions:

```
How many passes of UETP do you wish to run [1]?
How many simulated user loads do you want [4]?
Do you want Long or Short report format [Long]?
```

Press Return after each prompt. After you answer the last question, UETP initiates its entire sequence of tests, which run to completion without further input. The final message should look like the following:

```
*****
*                                     *
*   END OF UETP PASS 1 AT 22-JUN-1995 16:30:09.38   *
*                                     *
*****
```

Note

If you want to run UETP without using the default responses, refer to Section 17.4, which explains your options.

5. After UETP runs, check the log files for errors. If testing completes successfully, the OpenVMS operating system is in proper working order. If UETP does not complete successfully, refer to Section 17.6 for information on troubleshooting.

Note

After a run of UETP, you should run the Error Log utility to check for hardware problems that can occur during a run of UETP. For information on running the Error Log utility, refer to the *OpenVMS System Management Utilities Reference Manual*.

17.2 Preparing to Use UETP

This section contains detailed instructions for running UETP, including:

- Logging in
- Using the [SYSTEST] directory

Testing the System with UETP

17.2 Preparing to Use UETP

17.2.1 Logging In

Obtain the SYSTEST password from your system manager. Log in to the SYSTEST account from the console terminal as follows:

```
Username: SYSTEST
Password:
```

Note

Because SYSTEST has privileges, unauthorized use of this account can compromise the security of your system.

UETP will fail if you do not run the test from the SYSTEST account. Also, if you try to run UETP from a terminal other than the console terminal, the device test phase displays an error message stating that the terminal you are using is unavailable for testing. You can ignore this message.

After you log in to the SYSTEST account, enter the command `SHOW USERS` to make sure no user programs are running and no user volumes are mounted. UETP requires exclusive use of system resources. If you ignore this restriction, UETP can interfere with applications that depend on these resources.

Note

The information contained in Section 17.7.2 can help you identify and solve problems, including wrong quotas, privileges, or accounts, that could occur when you are running UETP. Refer to this section before you run UETP.

17.2.2 Using the SYSTEST Directories

If you logged in successfully, your default directory is [SYSTEST] on the system disk. UETP uses this directory to hold all the files used by UETP command procedure (UETP.COM) and temporary files used by UETP during testing.

On a typical system, the DCL command `SHOW LOGICAL` displays the translation of the logical name `SYS$TEST`:

```
$ SHOW LOGICAL SYS$TEST
"SYS$TEST" = "SYS$SYSROOT:[SYSTEST]" (LNM$SYSTEM_TABLE)
```

To use UETP to test a particular disk, such as a scratch disk, create either a [SYSTEST] directory or a [SYS0.SYSTEST] directory on that disk. Section 17.3.3 discusses setting up scratch disks for testing.

17.3 Setting Up the Devices to Be Tested

After you log in, set up the devices on the system for UETP testing, as described in the following sections. Note that your system might not have all the devices described in this section.

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

17.3.1 Check Your Devices

Examine all devices that UETP will use to be sure that the following conditions exist:

- All devices you want to test are turned on and are on line.
- Scratch disks are initialized and mounted.
- Disks contain a directory named [SYSTEST] with OWNER_UIC=[1,7]. Use the CREATE/DIRECTORY command if the [SYSTEST] directory does not exist on the disk.
- Scratch magnetic tape reels are *physically* mounted on each drive you want tested and are initialized with the label UETP (using the DCL command INITIALIZE). Make sure magnetic tape reels contain at least 600 feet of tape.
- Scratch tape cartridges have been inserted in each drive you want to test and are mounted and initialized with the label UETP (using the DCL command INITIALIZE).
- Line printers and hardcopy terminals have plenty of paper.
- Terminal characteristics and baud rate are set correctly (see the user's guide for your terminal).

Note that some communications devices discussed in this section must be set up by Digital Services.

17.3.2 System Disk Space Required

Before running UETP, be sure that the system disk has at least 1200 blocks available. Note that systems running more than 20 load test processes can require a minimum of 2000 available blocks. If you run multiple passes of UETP, log files will accumulate in the default directory and further reduce the amount of disk space available for subsequent passes.

If disk quotas are enabled on the system disk, disable them before you run UETP.

17.3.3 How UETP Works on Disks

The disk test phase of UETP uses most of the available free space on each testable disk in the following manner:

- On each testable disk, the device test phase tries to create two files. The size of these files depends on how much free space is available on the disk. Usually the test creates each file with 0.1% of the free space on the disk. However, if the disk is almost full, the test creates files that are 5 blocks. If the test cannot create 5 block files, it fails. Only the initial file creation can cause the device test to fail because it lacks disk space.
- The test randomly reads and writes blocks of data to the files. After every multiple of 20 writes for each file, the test tries to extend the file. The size of this extension is either 5% of the free disk space or 5 blocks if the file was created with 5 blocks. This process of extension continues until the combined space of the files reaches 75% of the free disk space.

By creating and extending fragmented files in this way, UETP exercises the disk. This allows the test to check for exceeded quotas or a full disk, and to adjust for the amount of available disk space.

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

As with other disks, shadow sets and volume sets can be tested with UETP; the expectation is that the individual members will be listed as untestable during UETINIDEV (initialization of UETP). UETINIDEV lists errors when testing using a shadow set during the system disk (UETDISK00) pass, however, the shadow set is listed as testable. When testing using a volume set, errors will be noted against all but relative volume number 1, and all but relative volume 1 will be listed as untestable at the end of UETINIDEV.

17.3.4 Prepare Disk Drives

To prepare each disk drive in the system for UETP testing, use the following procedure:

1. Place a scratch disk in the drive and spin up the drive. If a scratch disk is not available, use any disk with a substantial amount of free space; UETP does not overwrite existing files on any volume. If your scratch disk contains files that you want to keep, do not initialize the disk; go to Step 3.
2. If the disk does not contain files you want to save, initialize it. For example:

```
$ INITIALIZE DUA1: TEST1
```

This command initializes DUA1 and assigns the volume label TEST1 to the disk. All volumes must have unique labels.

3. Mount the disk. For example:

```
$ MOUNT/SYSTEM DUA1: TEST1
```

This command mounts the volume labeled TEST1 on DUA1. The /SYSTEM qualifier indicates that you are making the volume available to all users on the system.

4. UETP uses the [SYSTEST] directory when testing the disk. If the volume does not contain the directory [SYSTEST], you must create it. For example:

```
$ CREATE/DIRECTORY/OWNER_UIC=[1,7] DUA1:[SYSTEST]
```

This command creates a [SYSTEST] directory on DUA1 and assigns a user identification code (UIC) of [1,7]. The directory must have a UIC of [1,7] to run UETP.

If the disk you have mounted contains a root directory structure, you can create the [SYSTEST] directory in the [SYS0.] tree.

17.3.5 Magnetic Tape Drives

Set up magnetic tape drives that you want to test by performing the following steps:

1. Place a scratch magnetic tape with at least 600 feet of magnetic tape in the tape drive. Make sure that the write-enable ring is in place.
2. Position the magnetic tape at the BOT (beginning-of-tape) and put the drive on line.
3. Initialize each scratch magnetic tape with the label UETP. For example, if you have physically mounted a scratch magnetic tape on MUA1, enter the following command and press Return:

```
$ INITIALIZE MUA1: UETP
```

Magnetic tapes must be labeled UETP to be tested. As a safety feature, UETP does not test tapes that have been mounted with the MOUNT command.

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

If you encounter a problem initializing the magnetic tape or if the test has a problem accessing the magnetic tape, refer to the description of the INITIALIZE command in the *OpenVMS DCL Dictionary*.

17.3.6 Tape Cartridge Drives

To set up tape cartridge drives you want to test, do the following:

1. Insert a scratch tape cartridge in the tape cartridge drive.
2. Initialize the tape cartridge. For example:

```
$ INITIALIZE MUA0: UETP
```

Tape cartridges must be labeled UETP to be tested. As a safety feature, UETP does not test tape cartridges that have been mounted with the MOUNT command.

If you encounter a problem initializing the tape cartridge, or if the test has a problem accessing the tape cartridge, refer to the description of the DCL INITIALIZE command in the *OpenVMS DCL Dictionary*.

TLZ04 Tape Drives

During the initialization phase, UETP sets a time limit of 6 minutes for a TLZ04 unit to complete the UETTAPE00 test. If the device does not complete the UETTAPE00 test within the allotted time, UETP displays a message similar to the following:

```
-UETP-E-TEXT, UETTAPE00.EXE testing controller MKA was stopped ($DELPRC)
  at 16:23:23.07 because the time out period (UETP$INIT_TIMEOUT)
  expired or because it seemed hung or because UETINIT01 was aborted.
```

To increase the timeout value, enter a command similar to the following before running UETP:

```
$ DEFINE/GROUP UETP$INIT_TIMEOUT "0000 00:08:00.00"
```

This example defines the initialization timeout value to 8 minutes.

17.3.7 Compact Disc Drives

To run UETP on an RRD40 or RRD50 compact disc drive, you must first load the test disc that you received with your compact disc drive unit.

17.3.8 Optical Disk Drives

To run UETP on an RV60 drive, set up the RV64 optical disk-storage system, by doing the following:

1. Use the Jukebox Control Software (JCS) to load an optical disk in each of the RV60 drives. JCS is a layered product on the OpenVMS operating system that comes with the RV64 and is responsible for controlling the robot arm that loads and unloads the disks.
2. Initialize the optical disks with the label UETP, but do not mount them.

UETP tests all the RV60s present in the RV64 simultaneously. Unlike the tape tests, UETP does not reinitialize the optical disks at the end of the test.

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

17.3.9 Terminals and Line Printers

Terminals and line printers must be turned on and on line to be tested by UETP. Check that line printers and hardcopy terminals have enough paper. The amount of paper required depends on the number of UETP passes that you plan to execute. Each pass requires two pages for each line printer and hardcopy terminal.

Check that all terminals are set to the correct baud rate and are assigned appropriate characteristics. (See the user's guide for your terminal.)

Spooled devices and devices allocated to queues fail the initialization phase of UETP and are not tested.

17.3.10 Ethernet Adapters

Make sure that no other processes are sharing the Ethernet adapter device when you run UETP.

Note

UETP will not test your Ethernet adapter if DECnet for OpenVMS or some other application has the device allocated.

Because either DECnet for OpenVMS or the LAT terminal server can try to use the Ethernet adapter (a shareable device), you must shut down DECnet and the LAT terminal server before you run the device test phase, if you want to test the Ethernet adapter.

17.3.11 DR11-W Data Interface (VAX Only)

VAX

The DR11-W data interface uses an internal logical loopback mode that tests all features except that of module connectors, cables, and transceivers.

Caution

Only Digital Services personnel can set up the DR11-W data interface for UETP testing.

Because random external patterns are generated during this operation, the user device or other processor might need to be isolated from the DR11-W data interface being tested until the testing is completed.

To test the DR11-W data interface properly, the E105 switchpack must be set as follows:

Switch 1	Switch 2	Switch 3	Switch 4	Switch 5
Off	On	Off	Off	On

When UETP testing is completed, restore the DR11-W data interface to the proper operating configuration. ♦

17.3.12 DRV11-WA Data Interface (VAX Only)

VAX

The DRV11-WA data interface is a general-purpose, 16-bit, parallel, direct memory access (DMA) data interface.

Caution

Only Digital Services personnel can set up the DRV11-WA data interface for UETP testing.

To prepare the DRV11-WA driver on a MicroVAX computer for UETP testing, be sure the following conditions exist:

- The jumpers on the DRV11-WA board are set to W2, W3, and W6.
- A loopback cable is connected to the DRV11-WA board.
- The DRV11-WA board occupies slots 8 to 12. If the DRV11-WA is in another location, timeout errors can occur.

When UETP testing is completed, restore the DRV11-WA to the proper operating configuration. ♦

17.3.13 DR750 or DR780 (DR32 Interface) (VAX Only)

VAX

The DR32 (DR750 or DR780) device is an interface adapter that connects the internal memory bus of a VAX processor to a user-accessible bus called the DR32 device interconnect (DDI).

Caution

Only Digital Services personnel can set up the DR750 or DR780 for UETP testing.

To prepare the DR750 or the DR780 for UETP testing, use the following procedure:

1. Copy the DR780 microcode file, XF780.ULD, from the diagnostic medium to SYS\$SYSTEM. Use the procedure described in the documentation provided with the DR780 Microcode Kit.
2. Turn off the power to the DR780.
3. Make the following DR780 backplane jumper changes:
 - a. Remove the jumper from W7 and W8
 - b. Add a jumper from E04M1 to E04R1
 - c. Add a jumper from E04M2 to E04R2
4. Disconnect the DDI cable from the DR780. This cable is either a BC06V-nn cable, which can be disconnected, or a BC06R-nn cable, which requires that you remove its paddle card from the backplane of the DR780.
5. Restore power to the DR780.

When UETP testing is completed, restore the DR750 or the DR780 to the proper operating configuration. ♦

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

17.3.14 Second LPA11-K Device

VAX

If you have two LPA11-K devices, be sure that each is given a systemwide logical name in the SYS\$MANAGER:LPA11STRT.COM file. The logical name for the first LPA11-K device should be LPA11\$0, and the logical name for the second LPA11-K device should be LPA11\$1. ♦

17.3.15 Devices That Are Not Tested

UETP does not test the following devices; their status has no effect on UETP execution:

- Devices that require operator interaction (such as card readers)
- Software devices (such as the null device and local memory mailboxes)

UETP does not have specific tests for UDA, HSC, or CI devices; they are tested implicitly by the disk, magnetic tape, and DECnet for OpenVMS tests.

UETP also does not test the console terminal or console drives. If you boot the system, log in, and start UETP, you have shown that these devices can be used.

17.3.16 VMScLuster Testing

Before you run UETP in a VMScLuster environment, you should check the SYSTEST_CLIG account. The SYSTEST_CLIG account parallels SYSTEST except that it is dedicated to running the cluster-integration test. The requirements for the SYSTEST_CLIG account are as follows:

- The account should be present in the user authorization file, exactly as distributed by Digital on each system in your VMScLuster.

Note

The SYSTEST_CLIG account could have been disabled during the OpenVMS upgrade procedure. If it was disabled, you must reenale the SYSTEST_CLIG account and give it a null password before you run UETP.

To reenale the SYSTEST_CLIG account, enter the following commands:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN AUTHORIZE
UAF> MODIFY /FLAGS=NODISUSER /NOPASSWORD SYSTEST_CLIG
UAF> EXIT
```

Note

Digital recommends that you disable the SYSTEST_CLIG account after testing has completed.

To disable the SYSTEST_CLIG account, enter the following commands:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN AUTHORIZE
UAF> MODIFY /FLAGS=DISUSER SYSTEST_CLIG
UAF> EXIT
```

- The privileges and quotas of the SYSTEST_CLIG account must match those of the SYSTEST account.

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

UETP requires little additional preparation for the cluster-integration test phase beyond the requirements for other UETP test phases. The additional requirements for cluster integration testing are as follows:

1. Your system must be a member of a VMScluster. If it is not, UETP displays a message and does not attempt to run the test.
2. Your system must use the same deadlock detection interval as the other systems in the VMScluster. (The deadlock detection interval is set by the system parameter DEADLOCK_WAIT. It is normally not changed from the default value, which is 10 seconds.)
3. The files UETCLIG00.COM and UETCLIG00.EXE, located in SYS\$TEST, are necessary for each system included in the test.
4. DECnet for OpenVMS must be set up between the VMScluster nodes; UETP uses DECnet for OpenVMS to create a process on those nodes. All checks that the test makes depend on its ability to create the SYSTEST_CLIG processes and to communicate with them using DECnet for OpenVMS software.
5. All operator terminals (OPA0:) should accept broadcast messages. To set the BROADCAST characteristic, enter the following command:

```
$ SET TERM/BROADCAST/PERM OPA0:
```

Nodes on which the operator's terminal (OPA0) is set to the NO BROADCAST terminal characteristic will generate the following error message during the cluster test:

```
*****
* UETCLIG00master *
* Error count = 1 *
*****
-UETP-E-TEXT, 0 operator consoles timed out on the cluster test warning
and 1 operator console rejected it.
-UETP-E-TEXT, Status returned was,
"%SYSTEM-F-DEVOFFLINE, device is not in configuration or not
available"
```

6. There must be a [SYSTEST] or [SYS0.SYSTEST] directory on some disk available to the VMScluster for each node (both OpenVMS and HSC) in the cluster. The test uses the same directory as the UETP disk test to create a file on each cluster node and to see if some other OpenVMS node in the cluster can share access to that file. There must be one such directory per node; the test continues with the next cluster node once it has finished with a file.
7. By default, the UETP cluster phase selects three nodes from the running VMScluster for deadlock, disk, and file access testing. However, if you want all cluster nodes tested, enter the following command before invoking UETP:

```
$ DEFINE/GROUP UETP$CTMODE ALL
```

17.3.17 Testing a Small-Disk System

After you install the OpenVMS VAX operating system on a small system disk (for example, an RZ23L), you might not have the 1200 blocks of free disk space required to run UETP successfully. If you do not have 1200 free blocks on your system disk, use VMSTAILOR to remove some files from the system disk before you run UETP. For instructions on using VMSTAILOR, refer to the OpenVMS upgrade and installation manual for your system.

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

17.3.18 DECnet for OpenVMS Phase

The DECnet for OpenVMS phase of UETP uses more system resources than other tests. You can, however, minimize disruptions to other users by running the test on the least busy node.

By default, the file UETDNET00.COM specifies the node from which the DECnet test will be run. To run the DECnet test on a different node, enter the following command before you invoke UETP:

```
$ DEFINE/GROUP UETP$NODE_ADDRESS node_address
```

This command equates the group logical name UETP\$NODE_ADDRESS to the node address of the node in your area on which you want to run the DECnet phase of UETP.

For example:

```
$ DEFINE/GROUP UETP$NODE_ADDRESS 9.999
```

You can also run the DECnet for OpenVMS test on a different node by entering the following command before you invoke UETP:

```
$ DEFINE/GROUP UETP$NODE_NAME "node" "username password"
```

Note

When you use the logical name UETP\$NODE_ADDRESS, UETP tests only the first active circuit found by NCP (Network Control Program). Otherwise, UETP tests all active testable circuits.

When you run UETP, a router node attempts to establish a connection between your node and the node defined by UETP\$NODE_ADDRESS or UETP\$NODE_NAME. Occasionally, the connection between your node and the router node can be busy or nonexistent. When this happens, the system displays the following error messages:

```
%NCP-F-CONNEC, Unable to connect to listener
-SYSTEM-F-REMRSRC, resources at the remote node were insufficient

%NCP-F-CONNEC, Unable to connect to listener
-SYSTEM-F-NOSUCHNODE, remote node is unknown
```

17.3.19 Vector Processors and the VVIEF (VAX Only)

VAX

UETP automatically loads all installed and enabled vector processors during the load phase, and automatically tests all installed and enabled vector processors during the device test phase.

If vector processors are available on the system, check for the VP number by entering the following commands:

```
$ x = F$GETSYI ("VP_NUMBER")
$ SHOW SYMBOL x
```

Multiply the value of *x* by 3. If the result is greater than the account PRCLM value, then you must increase the SYSTEST account PRCLM quota to match the returned result. For more information see Chapter 26.

Testing the System with UETP

17.3 Setting Up the Devices to Be Tested

However, UETP cannot load the VAX Vector Instruction Emulation facility (VVIEF) during the load phase, and will not automatically test VVIEF. To test VVIEF, you must do the following before running UETP:

1. Edit the file UETCONT00.DAT to add the following line:

```
Y Y UETVECTOR.EXE "DEVICE_TEST"
```

2. Make sure VVIEF was activated when the system was booted. To determine if the VVIEF was activated, enter the following DCL commands:

```
$ X = F$GETSYI("VECTOR_EMULATOR")
$ SHOW SYMBOL X
```

If the system displays a value of 1, VVIEF is loaded; if the system displays a value of 0, VVIEF is not loaded.

The VVIEF test can be executed as an individual test using the RUN command, as described in Section 17.8.2. ♦

17.4 Starting UETP

When you have logged in and prepared the system and devices, you are ready to begin the test.

To start UETP, enter the following command and press Return:

```
$ @UETP
```

UETP displays the following prompt:

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?
```

Throughout the startup dialog, brackets indicate the default value, which you can choose by pressing Return.

When running UETP for the first time, it is recommended that you choose the default value (ALL) and run all the phases. If you choose ALL, UETP displays three more questions, which are described in Section 17.4.2 through Section 17.4.4. If you want to run all the test phases, skip the next section.

17.4.1 Running a Subset of Phases

You can run a single phase by entering SUBSET or S in response to the following prompt:

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?
```

If you enter S or SUBSET, UETP prompts you for the phase you want to run as follows:

You can choose one or more of the following phases:

```
DEVICE, LOAD, DECNET, CLUSTER
```

Phases(s):

There is no default; enter one or more phase names from the list. Separate two or more phases with spaces or commas.

If your choice includes the LOAD phase, UETP displays three prompts:

```
How many passes of UETP do you wish to run [1]?
How many simulated user loads do you want [n]?
Do you want Long or Short report format [Long]?
```

If you exclude the LOAD phase from your list of choices, UETP responds with only two prompts: the first and the third.

Testing the System with UETP

17.4 Starting UETP

The next three sections discuss how you can respond to these questions. After you have answered the questions, the phase you have selected runs to completion.

17.4.2 Single Run Versus Multiple Passes

If you specified the default ALL or a subset of phases at the last prompt, UETP displays the following message:

How many passes of UETP do you wish to run [1]?

You can repeat the test run as many times as you want. If you enter 1 in response to the prompt (or press Return for the default), UETP stops after completing a single run. If you specify a number greater than 1, UETP restarts itself until it completes the specified number of passes.

You can run UETP once to check that the system is working, or many times to evaluate the system's response to continuous use. For example, a service technician who is interested only in verifying that a newly installed system works might run UETP once or twice. A manufacturing technician might let the system run for several hours as part of the system integration and test.

When you specify multiple UETP runs, you can request a short console log. (See Section 17.4.4.) Ensure that all line printers and hardcopy terminals have enough paper because each run requires two pages.

17.4.3 Defining User Load for Load Test

After you specify the number of passes, UETP prompts you as follows:

How many simulated user loads do you want [n]?

Note

UETP displays this prompt only if you choose to run the LOAD phase, either implicitly (by running all phases) or explicitly (by running a subset and specifying the LOAD phase).

The load test simulates a situation in which a number of users (detached processes) are competing for system resources. In response to this prompt, enter the number of users you want to simulate for this test. The number in brackets is the default value that UETP computed for your system. The default value depends on the amount of memory and the paging and swapping space that your system has allocated.

Although the given default value is the best choice, you can increase or decrease the user load by entering your own response to the prompt. However, be aware that an increase can cause the test to fail because of insufficient resources.

If you want to see UETP display the user-load equation as it runs, see Section 17.6.2.

17.4.4 Report Formats

The following prompt allows you to choose between long or short report formats:

Do you want Long or Short report format [Long]?

17.4.4.1 Long Report Format

If you choose the long report format (the default), UETP sends the following information to the console terminal:

- All error messages
- All output generated at the beginning of all phases and tests
- All output generated at the end of all phases and tests

UETP records all its output in the UETP.LOG file, regardless of your response to this question.

In many cases, it might not be convenient to have UETP write the bulk of its output to the terminal. For example, if you run UETP from a hardcopy terminal, the output printing can slow the progress of the tests. This delay might not be a problem if you have requested only one run; however, you might prefer to use the short format if you intend to run multiple passes of UETP from a hardcopy terminal.

17.4.4.2 Short Report Format

If you request the short format, UETP displays status information at the console, such as error messages and notifications of the beginning and end of each phase. This information enables you to determine whether UETP is proceeding normally. If the short console log indicates a problem, you can look at the file UETP.LOG for further information. UETP.LOG contains all the output generated by the various phases, as well as the status information displayed at the console.

After you choose the report format, UETP initiates its sequence of tests and runs to completion. If UETP does not complete successfully, refer to Section 17.6 for troubleshooting information.

17.5 Stopping a UETP Operation

At the end of a UETP pass, the master command procedure UETP.COM displays the time at which the pass ended. In addition, UETP.COM determines whether UETP needs to be restarted. You can request multiple passes when you start up the test package. (See Section 17.4.2.)

At the end of an entire UETP run, UETP.COM deletes temporary files and does other cleanup activities.

Pressing Ctrl/Y or Ctrl/C lets you terminate a UETP run before it completes normally. Normal completion of a UETP run, however, includes the deletion of miscellaneous files that have been created by UETP for the purpose of testing. Using Ctrl/Y or Ctrl/C can interrupt or prevent these cleanup procedures.

The effect of these control characters depends on what part of UETP you are executing. For an explanation of the organization of UETP and its components, refer to Section 17.8.

17.5.1 Using Ctrl/Y

Press Ctrl/Y to abort a UETP run. Note, however, that cleanup of files and network processes in the [SYSTEST] directory might not be complete.

If you are running an individual test image, pressing Ctrl/Y interrupts the current UETP test and temporarily returns control to the command interpreter. While the test is interrupted, you can enter a subset of DCL commands that are executed within the command interpreter and do not cause the current image to exit.

Testing the System with UETP

17.5 Stopping a UETP Operation

17.5.2 Using DCL Commands

The *OpenVMS User's Manual* contains a table of commands that you can use within the command interpreter. In addition, you can enter any of the following commands:

- The CONTINUE command continues the test from the point of interruption (except during execution of the cluster test).
- The STOP command terminates the test; the test aborts and control returns to the command interpreter.

Note

Using the STOP command can prevent cleanup procedures from executing normally. You should use the EXIT command if you want the image to do cleanup procedures before terminating.

- The EXIT command executes cleanup procedures and terminates the test (except during execution of the cluster test); control returns to the command interpreter.

If you enter any DCL command other than those that execute within the command interpreter, the test does cleanup procedures and terminates, and the DCL command executes.

17.5.3 Using Ctrl/C

Press Ctrl/C to interrupt a UETP run. You cannot continue the same test phase after you press Ctrl/C. UETP automatically goes to the next phase in the master command procedure.

Some UETP phases react to Ctrl/C by cleaning up all activity and terminating immediately. These tests display the following message when they are started:

```
%UETP-I-ABORTC, 'testname' to abort this test, type ^C
```

The phases that do not display the previous message terminate all processes they have started. These processes might not have a chance to complete normal cleanup procedures.

If you are running an individual test image, however, you can use Ctrl/C to terminate the execution of the image and complete cleanup procedures.

Note that Ctrl/C does not complete cleanup procedures for the cluster test.

17.6 Troubleshooting: An Overview

This section explains the role of UETP in interpreting operational errors in an OpenVMS operating system. See Section 17.7 for a discussion of common errors that can appear in a UETP run and describes how to correct them.

17.6.1 Error Logging and Diagnostics

When UETP encounters an error, it reacts like a user program. It either returns an error message and continues, or it reports a fatal error and terminates the image or phase. In either case, UETP assumes the hardware is operating properly and it does not attempt to diagnose the error.

Testing the System with UETP

17.6 Troubleshooting: An Overview

If the cause of an error is not readily apparent, use the following methods to diagnose the error:

- OpenVMS Error Log utility (ERROR LOG)—Run ERROR LOG to obtain a detailed report of hardware and system errors. ERROR LOG reports provide information about the state of the hardware device and I/O request at the time of each error. For information about running ERROR LOG refer to the *OpenVMS System Management Utilities Reference Manual*.
- Diagnostic facilities—Use the diagnostic facilities to test exhaustively a device or medium to isolate the source of the error.

17.6.2 Interpreting UETP Output

You can monitor the progress of UETP tests at the terminal from which they were started. This terminal always displays status information, such as messages that announce the beginning and end of each phase and messages that signal an error.

The tests send other types of output to various log files, depending on how you started the tests. (See Section 17.6.7.) The log files contain output generated by the test procedures. Even if UETP completes successfully, with no errors displayed at the terminal, it is good practice to check these log files for errors. Furthermore, when errors are displayed at the terminal, check the log files for more information about their origin and nature.

Each test returns a final completion status to the test controller image, UETPHAS00, using a termination mailbox. This completion status is an unsigned longword integer denoting a condition value. As a troubleshooting aid, UETPHAS00 displays the test's final completion status using the \$FAO and \$GETMSG system services.

Sometimes, however, the \$FAO service needs additional information that cannot be provided using the termination mailbox. When this happens, UETP displays an error message similar to the following:

```
UETP-E-ABORT, !AS aborted at !%D
```

When UETP displays these types of error messages, check the log files for more information. You can also run the individual test to attempt to diagnose the problem.

The error messages that appear at the terminal and within the log files have two basic sources:

- UETP tests
- System components that are tested

If you need help interpreting the messages, use the OpenVMS Help Message utility (Help Message) or refer either to the *OpenVMS System Messages and Recovery Procedures Reference Manual* or to the manual that describes the individual system component.

17.6.3 Displaying Information on Your Screen

Several parts of UETP, such as some device tests, UETINIT00.EXE, UETCLIG00.EXE, and UETDNET00.COM, let you obtain additional information concerning the progress of the test run or the problems the test encounters. Because this information is usually insignificant, it is not displayed on the screen.

Testing the System with UETP

17.6 Troubleshooting: An Overview

To view the information, enter the following command to define the logical name MODE and run the program:

```
$ DEFINE MODE DUMP
```

17.6.4 Example Screen Display (VAX Only)

VAX

The following example shows the output for UETINIT00.EXE on a MicroVAX 3600 system:

```
$ RUN UETINIT00
```

```
Welcome to VAX/VMS UETP Version 6.2
```

```
%UETP-I-ABORTC, UETINIT00 to abort this test, type ^C
```

```
You are running on a MicroVAX 3600 Series CPU with 65536 pages of memory.  
The system was booted from _DUA0:[SYS0.].
```

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?  
How many passes of UETP do you wish to run [1]?
```

```
The default number of loads is the minimum result of
```

```
1) CPU_SCALE * ((MEM_FREE + MEM_MODIFY) / (WS_SIZE * PER_WS_INUSE))  
   2.50 * (( 28126 + 312) / ( 1024 * 0.20)) = 347  
2) Free process slots = 197  
3) Free page file pages / Typical use of page file pages per process  
   96920 / 1000 = 96
```

```
How many simulated user loads do you want [96]?  
Do you want Long or Short report format [Long]?
```

```
UETP starting at 22-JUN-1995 09:08:26.71 with parameters:  
DEVICE LOAD DECNET CLUSTER phases, 1 pass, 96 loads, long report.  
$
```

This program does not initiate any phase; it displays the equation used by UETP to determine user load and the specific factors that are employed in the current run.

Respond to the questions by pressing Return. After you respond to the first prompt, the program displays the expressions that determine the default number of simultaneous processes. The following definitions apply:

- CPU_SCALE refers to the relative processing power of the CPU in relation to a VAX 11/780 computer. For example, a MicroVAX 3600 computer has a CPU_SCALE of 2.5 because it has 2.5 times the processing power of a VAX 11/780 (1.0) computer.
- MEM_FREE represents memory in pages available to users.
- MEM_MODIFY represents memory pages on the modified page list.
- WS_SIZE represents working set size.
- PER_WS_INUSE represents typical percentage of the working set in active use for each process.

UETINIT00 also displays the specific values represented by the expressions. In this example, UETP selects 96 as the default for simulated user loads, because 96 is the minimum result of the three expressions.

You should deassign the logical name MODE before running UETP, unless you prefer to see the previous breakdown every time you run UETP. ♦

17.6.5 Example Screen Display (Alpha Only)

Alpha

The following example shows the output for UETINIT00.EXE on an Alpha system:

```
$ RUN UETINIT00.EXE
```

```
Welcome to OpenVMS Alpha UETP Version 6.2
```

```
%UETP-I-ABORTC, UETINIT00 to abort this test, type ^C
```

```
You are running on a DEC 4000 Model 610 CPU.
```

```
The system was booted from _COB3$DKA0:[SYS0.].
```

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?
```

```
How many passes of UETP do you wish to run [1]?
```

```
The default number of loads is the minimum result of
```

```
1) (MEM_FREE + MEM_MODIFY) / ( WS_SIZE )  
   ( 215696 +      11136 ) / (   4000 )                = 56  
2) Free process slots                                = 281  
3) Free page file pages / Typical use of blocks per process  
   199936 / 1000 = 199
```

```
How many simulated user loads do you want [56]?
```

```
Do you want Long or Short report format [Long]?
```

```
UETP starting at 22-APR-1995 12:20:01.32 with parameters:
```

```
DEVICE LOAD DECNET CLUSTER phases, 1 pass, 56 loads, long report.
```

```
$
```

This program does not initiate any phase; it displays the equation used by UETP to determine user load and the specific factors that are employed in the current run.

Respond to the questions by pressing the Return key. After you respond to the first prompt, the program displays the expressions that determine the default number of simultaneous processes. The following definitions apply:

- MEM_FREE represents memory in pagelets available to users.
- MEM_MODIFY represents memory pagelets on the modified page list.
- WS_SIZE represents working set size in pagelets.

UETINIT00 also displays the specific values represented by the expressions. In this example, UETP selects 56 as the default for simulated user loads, because 56 is the minimum result of the three expressions.

You should deassign the logical name MODE before running UETP, unless you prefer to see the previous breakdown every time you run UETP. ♦

17.6.6 Defining a Remote Node for UETP Ethernet Testing

Occasionally during the UETUNAS00 test, it is difficult to determine whether the problem reports concern the device under test or the remote device. The easiest way to ensure proper error reporting is to define a *good turnaround*. A good turnaround is a remote node that you know turns around Ethernet packets correctly and is up and waiting in the ready state.

You can make the UETUNAS00 test use a known good turnaround by performing the following actions. In the commands that follow, assume that the *good* device is on node BETA and that node BETA is already defined in the network database.

Testing the System with UETP

17.6 Troubleshooting: An Overview

1. Find the address of the good Ethernet node by using the Network Control Program (NCP). To use NCP, the following conditions must apply:
 - DECnet for OpenVMS must be up and running on the system.
 - The account you are using must have TMPMBX and NETMBX privileges.

Enter the following commands and press Return:

```
$ RUN SYS$SYSTEM:NCP
NCP> TELL BETA SHOW EXECUTOR STATUS
```

If node BETA has not been defined in your network database, NCP displays an error message. In this event, specify another good node and retry the command. Otherwise, see your system or network manager.

NCP displays information similar to the following:

```
Node Volatile Status as of 22-JUN-1995 16:13:02
Executor node = 19.007 (BETA)

State           = on
Physical address = AA-00-03-00-76-D3
Active links    = 6
Delay           = 1
```

2. Use the displayed *physical address* (in this case, AA00030076D3) to define the logical name TESTNIADR to point to the good turnaround. Note that you do not specify the hyphens (-).

First, log in to the SYSTEST account. Then enter the following command:

```
$ DEFINE/SYSTEM TESTNIADR AA00030076D3
```

3. Run UETP.
4. When UETP has completed, deassign the logical name TESTNIADR by entering the following command:

```
$ DEASSIGN/SYSTEM TESTNIADR
```

17.6.7 Log Files

UETP stores all information generated by all UETP tests and phases from its current run in one or more UETP.LOG files, and it stores the information from the previous run in one or more OLDUETP.LOG files. If a run of UETP involves multiple passes, there will be one UETP.LOG or one OLDUETP.LOG file for each pass.

At the beginning of a run, UETP deletes all OLDUETP.LOG files, and renames any UETP.LOG files to equivalent versions of OLDUETP.LOG. Then UETP creates a new UETP.LOG file and stores the information from the current pass in it. Subsequent passes of UETP create higher versions of UETP.LOG. Therefore, at the end of a run of UETP that involves multiple passes, there is one UETP.LOG file for each pass. In producing the files UETP.LOG and OLDUETP.LOG, UETP provides the output from the two most recent runs.

The cluster test creates a NETSERVER.LOG file in SYS\$TEST for each pass on each system included in the run. If the test is unable to report errors (for example, if the connection to another node is lost), the NETSERVER.LOG file on that node contains the result of the test run on that node. UETP does not purge or delete NETSERVER.LOG files; therefore, you must delete them occasionally to recover disk space.

If a UETP run does not complete normally, SYS\$TEST can contain other log files. Ordinarily these log files are concatenated and placed within UETP.LOG. You can use any log files that appear on the system disk for error checking, but you must delete these log files before you run any new tests. You can delete these log files yourself or rerun the entire UETP, which checks for old UETP.LOG files and deletes them.

17.7 Troubleshooting: Possible UETP Errors

This section is intended to help you identify and solve problems you can encounter running UETP. You should refer to this section if you need help understanding a system failure and isolating its cause. This section is not intended as a repair manual and is not expected to diagnose any flaws in your system. It should, however, help you to interpret and act upon the information in the error messages.

If you are unable to correct an error after following the steps in this section, you should contact your Digital Services representative. Any information you can supply about the measures you have taken to isolate the problem will help your Digital Services representative diagnose the problem.

17.7.1 Summary of Common Failures

The following are the most common failures encountered while running UETP:

- Wrong quotas, privileges, or account
- UETINIT01 failure
- UETVECTOR failure (VAX computers only)
- Ethernet device allocated or in use by another application
- Insufficient disk space
- Incorrect VMSccluster setup
- Problems during the load test
- DECnet for OpenVMS error
- Errors logged but not displayed
- No process control block (PCB) or swap slots
- System hangups
- Lack of default access for the file access listener (FAL) object
- Bugchecks and machine checks

The sections that follow describe these errors and offer the best course of action for dealing with each one.

17.7.2 Wrong Quotas, Privileges, or Account

If your assigned quotas or privileges do not match standard quotas and privileges for the SYSTEST account, UETP displays the following error message:

```
*****
* UETINIT00
* Error count = 1 *
*****
-UETP-W-TEXT, The following:
```


Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

OPER privilege,
BIOLM quota,
ENQLM quota,
FILLM quota,

are nonstandard for the SYSTEST account and may result in UETP errors.

This message informs you that the OPER privilege and the BIOLM, ENQLM, and FILLM quotas either are not assigned correctly or are not assigned at all.

Note

UETP displays a similar message if you run the cluster integration test phase and the privileges and quotas for the SYSTEST_CLIG account are incorrect. The SYSTEST and SYSTEST_CLIG accounts require the same privileges and quotas. Take the action described in this section for both accounts.

Solution

To correct the problem, use the following procedure:

1. Display all privileges and quotas in effect for the SYSTEST account using the Authorize utility (AUTHORIZE) as follows:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYS$SYSTEM:AUTHORIZE
UAF> SHOW SYSTEST

Username: SYSTEST                      Owner: SYSTEST-UETP
Account:  SYSTEST                      UIC:    [1,7] ([SYSTEST])
CLI:      DCL                          Tables: DCLTABLES
Default:  SYS$SYSROOT:[SYSTEST]
LGICMD:   LOGIN
Login Flags:
Primary days: Mon Tue Wed Thu Fri Sat Sun
Secondary days:
No access restrictions
Expiration:          (none)          Pwdminimum: 8      Login Fails: 0
Pwdlifetime:         14 00:00        Pwdchange: 22-JUN-1995 10:12
Last Login:          (none) (interactive), (none) (non-interactive)
Maxjobs:             0 Fillm:         100 Bytln:         65536
Maxacctjobs:         0 Shrfillm:      0 Pbytln:          0
Maxdetach:           0 BIOLm:         12 JTquota:        1024
Prclm:               12 DIOLm:        55 WSdef:          256
Prio:                4 ASTlm:         100 WSquo:          512
Queprio:             0 TQElm:         20 WSextent:       2048
CPU:                 (none) Enqlm:     300 Pgflquo:       20480

Authorized Privileges:
CMKRNL CMEXEC SYSNAM GRPNAM DETACH DIAGNOSE LOG_IO GROUP
PRMCEB PRMMBX SETPRV TMPMBX NETMBX VOLPRO PHY_IO SYSPRV
Default Privileges:
CMKRNL CMEXEC SYSNAM GRPNAM DETACH DIAGNOSE LOG_IO GROUP
PRMCEB PRMMBX SETPRV TMPMBX NETMBX VOLPRO PHY_IO SYSPRV
UAF> SHOW SYSTEST_CLIG
.
.
.
UAF> EXIT
```

2. Make sure the default privileges and quotas assigned to the account match the following:

Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

Privileges

CMKRNL	CMEXEC	NETMBX	DIAGNOSE
DETACH	PRMCEB	PRMMBX	PHY_IO
GRPNAM	TMPMBX	VOLPRO	LOG_IO
SYSNAM	SYSRV	SETPRV	GROUP

Quotas

BIOLM: 18	PRCLM: 12
DIOLM: 55	ASTLM: 100
FILLM: 100	BYTLM: 65536
TQELM: 20	CPU: no limit
ENQLM: 300	PGFLQUOTA: 20480
WSDEFAULT: 256	WSQUOTA: 512
WSEXTENT: 2048	

3. If any privileges or quotas are incorrect, run **AUTHORIZE** to correct them.

If you are logged in to the wrong account, the following error message asks you to log in to the **SYSTEST** account:

```
$ @UETP
*****
* UETINIT00 *
* Error count = 1 *
*****
-UETP-E-ABORT, UETINIT00 aborted at 22-JUN-1995 14:24:10.13
-UETP-E-TEXT, You are logged in to the wrong account.
Please log in to the SYSTEST account.
$
```

You must run UETP from the **SYSTEST** account.

17.7.3 UETINIT01 Failure

UETINIT01 failures are related to peripheral devices; this type of error message can indicate any of the following:

- Device failure
- Device not supported or not mounted
- Device allocated to another user
- Device write locked
- Lost vacuum on a magnetic tape drive
- Drive off line

In some cases, the corrective action is specified explicitly in the error message. For example, you can receive a message from the operator communication manager (OPCOM) informing you of a problem and recommending a corrective measure:

```
%OPCOM, 22-JUN-1995 14:10:52.96, request 1, from user SYSTEST
Please mount volume UETP in device _MTA0:
%MOUNT-I-OPRQST, Please mount volume UETP in device _MTA0:
```


Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

Other error messages can relate information in which the solution is specified implicitly:

```
%UETP-S-BEGIN, UETDISK00 beginning at 22-JUN-1995 13:34:46.03
```

```
*****
```

```
*   DISK_DRA   *
```

```
*   Error count = 1   *
```

```
*****
```

```
-UETP-E-TEXT, RMS file error in file DRA0:DRA00.TST
```

```
-RMS-E-DNR, device not ready or not mounted
```

```
%UETP-S-ENDED, UETDISK00 ended at 22-JUN-1995 13:34:46.80
```

This message tells you that a disk drive is either not ready or not mounted. From this information, you know where to look for the cause of the failure (at the disk drive). If you cannot see the cause of the problem immediately, check the setup instructions in Section 17.3.

In other cases, the cause of a failure might not be obvious from the information in the message. The problem can be related to hardware rather than software. For example, the Ethernet adapter test may produce one of the following messages if UETP does not have exclusive access to the Ethernet adapter:

- Intermodule cable unplugged
- Self-test failure code 0000000

To run the self-test diagnostic on the Ethernet adapter successfully, UETP needs exclusive access to the adapter. As explained in Section 17.3.10, you must shut down DECnet and the LAT terminal server before running the UETP device test phase if you want to test the Ethernet adapter.

Solution

To determine where or when the failure occurs in the execution of UETP, use the following procedure:

- Run the device test individually. (See Section 17.4.1.) By doing this, you can determine if the failure can be re-created, and you can isolate the cause of the problem by reproducing it using the least amount of software possible.

For example, if the failure occurs only when you run the entire device phase, and not when you run the affected device test individually, you can conclude the problem is related to device interaction. Conversely, if you can re-create the error by running the single device test, then you have proved that the error is not related to device interaction.

- Run the device test with different media. If your run of the single device test succeeded in reproducing the error, the magnetic tape or disk media could be defective. Running the same test with different media determines whether the original media caused the problem.
- Call Digital Services. If you have tried all the previous steps without solving the problem, you should contact your Digital Services representative.

17.7.4 UETVECTOR Failure (VAX Only)

VAX

UETP displays a message similar to the following to signal a vector processor failure:

```
*****
*   UETVECTOR           *
*   Error count = 1     *
*****
%PPL-S-CREATED_SOME, created some of those requested - partial success
-UETP-E-SUBSPNERR, Error spawning subordinate process.
-UETP-E-SCHCTXERR, Error scheduling vector context test subprocess.
-UETP-E-VECCTXERR, Error encountered during vector context testing.
%UETP-I-ENDED, UETVECTOR_0000 ended at 22-JUN-1995 07:37:00.59
```

Solution

See Section 17.3.19 for the correct setup for vector processor testing. ♦

17.7.5 Device Allocated or in Use by Another Application

If DECnet for OpenVMS software or the LAT software is running during the DEVICE phase, the UETUNAS00 test displays the following message:

```
-UETP-W-TEXT, Device is in use by DECnet or another application
```

Other UETP communication device tests display the following message:

```
SYSTEM-W-DEVALLOC, device already allocated to another user
```

Solution

If you want to run the device test on the Ethernet Adapter, shut down DECnet and LAT software before beginning the test.

17.7.6 Insufficient Disk Space

When you run continuous passes of UETP, log files accumulate on the disk from which UETP was run. These files reduce the amount of free disk space available for each successive pass. If the amount of disk space available becomes too small for the current load, the following error message appears:

```
%UETP-S-BEGIN, UETDISK00 beginning at 22-JUN-1995 08:12:24.34
%UETP-I-ABORTC, DISK_DJA to abort this test, type ^C

*****
*   DISK_DJA           *
*   Error count = 1     *
*****
-UETP-F-TEXT, RMS file error in file DJA0:DJA00.TST
-RMS-F-FUL, device full (insufficient space for allocation)

*****
*   DISK_DJA           *
*   Error count = 2     *
*****
-UETP-F-TEXT, RMS file error in file DJA0:DJA01.TST
-RMS-F-FUL, device full (insufficient space for allocation)
%UETP-E-DESTP, DISK_DJA stopped testing DJA unit 0 at 08:12:36.91
%UETP-S-ENDED, UETDISK00 ended at 22-JUN-1995 08:12:37.98
```

Solution

Make more space available on the disk. You can do this by using one or more of the following techniques:

- Delete unnecessary files to create more space.
- Purge files, if multiple versions exist.
- Mount a volume with sufficient space.

Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

- Check for disk quotas that might be enabled on the disk. If disk quotas are enabled, either disable or increase them. (See the *OpenVMS System Management Utilities Reference Manual* for a description of the Disk Quota utility.)
- Run VMSTAILOR if you have a small-disk system. See the upgrade and installation manual for your operating system for more information.

See Section 17.2.2 and Section 17.3.3 for a further discussion of disk space.

17.7.7 Incorrect Setup of a VMSccluster System

Most problems that can occur during the cluster-integration test are related to improper setup of the VMSccluster system or of UETP on the VMSccluster. These problems are most likely to occur at the following stages of the VMSccluster test:

- Near the beginning, when processes on OpenVMS nodes are started
- Toward the end, when cluster file access is checked

The cluster test phase shows that various OpenVMS nodes in your cluster can simultaneously access files on selected nodes in the cluster. First, UETP tries to create a file on a disk drive that is accessible to the other selected nodes in the cluster. The following are the requirements for creating a file in the cluster test phase:

- A [SYSTEST] directory must exist on the disk in either the master file directory (MFD) or in the root directory [SYS0.].
- The protection for [SYSTEST] directory must be set to allow the SYSTEST account to create a file in it.

If UETP is unable to find a suitable device on a certain node, the test displays a warning message and proceeds to the next cluster node.

Nodes on which the operator's terminal (OPA0) is set to the NO BROADCAST terminal characteristic will generate the following error message during the cluster test:

```
*****
* UETCLIG00master *
* Error count = 1 *
*****
-UETP-E-TEXT, 0 operator consoles timed out on the cluster test warning
and 1 operator console rejected it.
-UETP-E-TEXT, Status returned was,
"%SYSTEM-F-DEVOFFLINE, device is not in configuration or not
available"
```

Disregard this message if OPA0 is set to NO BROADCAST.

Solution

Whenever you suspect a problem, examine the SYS\$TEST:NETSERVER.LOG file that was created when the SYSTEST_CLIG process was created. This file can contain additional error information that could not be transmitted to the node running the test. If it was not possible to create the SYSTEST_CLIG process on some node, the system accounting file for that node might contain a final process status in a process termination record.

The following problems can occur during a cluster test:

- Logging in at other nodes—This problem is due to incorrect setup for the cluster test at the remote OpenVMS node. For example, if you specified a

Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

password for the SYSTEST_CLIG account or if you disabled the SYSTEST_CLIG account, the test displays the following message:

```
%SYSTEM-F-INVLOGIN, login information invalid at remote node
```

Refer to Section 17.3.16 and Section 17.6.6 for information on preparing for VMScluster testing.

- Communicating with other nodes—A message indicates a DECnet problem. Check the NETSERVER.LOG file on the affected node to determine the cause.
- Taking out locks or detecting deadlocks—The most likely cause of this problem is that you are not logged in to the SYSTEST account. Another possibility is that your cluster is not configured properly.
- Creating files on VMScluster nodes—This problem is due to incorrect setup for the cluster test; refer to Section 17.3.16 for information about preparing for VMScluster testing.

17.7.8 Problems During the Load Test

A variety of errors can occur during the load test because the command procedures that are started during the tests run several utilities and do many functions. Tracking a problem can be difficult because UETP deletes the log files that are generated during the load test. (See Section 17.8.3.)

Solution

If a problem occurs during the load test and the cause is not obvious, you can modify UETP.COM to preserve the log files as follows:

1. Add the /NODELETE qualifier to the following line:

```
$ TCNTRL UETLOAD00.DAT/PARALLEL_COUNT='LOADS/REPORT_TYPE='REPORT
```

2. Delete or comment out the following line:

```
$ DELETE UETLO*.LOG;*
```

Rerun the load test with these changes to try to re-create the problem.

If you re-create the problem, look at the contents of the appropriate log file. You can determine which log file to read by understanding the scheme by which the load test names its processes and log files. (The log file names are derived from the process names.)

The load test creates processes that are named in the following format:

UETLOADnn_nnnn

For example:

```
%UETP-I-BEGIN, UETLOAD00 beginning at 22-JUN-1995 15:45:08.97
%UETP-I-BEGIN, UETLOAD02_0000 beginning at 22-JUN-1995 15:45:09.42
%UETP-I-BEGIN, UETLOAD03_0001 beginning at 22-JUN-1995 15:45:09.63
%UETP-I-BEGIN, UETLOAD04_0002 beginning at 22-JUN-1995 15:45:10.76
%UETP-I-BEGIN, UETLOAD05_0003 beginning at 22-JUN-1995 15:45:11.28
%UETP-I-BEGIN, UETLOAD06_0004 beginning at 22-JUN-1995 15:45:12.56
%UETP-I-BEGIN, UETLOAD07_0005 beginning at 22-JUN-1995 15:45:13.81
%UETP-I-BEGIN, UETLOAD08_0006 beginning at 22-JUN-1995 15:45:14.95
%UETP-I-BEGIN, UETLOAD09_0007 beginning at 22-JUN-1995 15:45:16.99
%UETP-I-BEGIN, UETLOAD10_0008 beginning at 22-JUN-1995 15:45:19.32
%UETP-I-BEGIN, UETLOAD11_0009 beginning at 22-JUN-1995 15:45:19.95
%UETP-I-BEGIN, UETLOAD02_0010 beginning at 22-JUN-1995 15:45:20.20
%UETP-I-BEGIN, UETLOAD03_0011 beginning at 22-JUN-1995 15:45:21.95
%UETP-I-BEGIN, UETLOAD04_0012 beginning at 22-JUN-1995 15:45:22.99
```


Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

Note that if more than 10 processes are created, the numbering sequence for the UETLOAD nn portion of the process name starts over at UETLOAD02; however, the 4 digits of the $nnnn$ portion continue to increase.

Each load test process creates two log files. The first log file is created by the test controller; the second log file is created by the process itself. The log file to look at for error information on any given load test process is the one that was created by the test controller (the first log file).

The load test log file derives its file name from the process name, appending the last four digits of the process name (from the $nnnn$ portion) to UETLO. The test-controller log file and the process log file for each process use the same file name; however, the process log file has the higher version number of the two. For example, the log files created by the process UETLOAD05_0003 would be named as follows:

UETLO0003.LOG;1 (test-controller log file)

UETLO0003.LOG;2 (process log file)

Make sure that you look at the log file with the lower version number; that file contains the load test commands and error information.

After you have isolated the problem, restore UETP.COM to its original state and delete the log files from the load test (UETLO*.LOG;*); failure to delete these files can result in disk space problems.

17.7.9 DECnet for OpenVMS Error

A DECnet error message can indicate that the network is unavailable.

Solution

- If DECnet for OpenVMS software is included in your system, determine whether the product authorization key (PAK) is registered by entering the following command:

```
$ SHOW LICENSE
```

If the PAK is not registered, invoke the License utility to register it by entering the following command:

```
$ @SYS$UPDATE:VMSLICENSE
```

For information about registering licenses, see the following:

- The OpenVMS Upgrade and Installation Manual for your operating system
- The *OpenVMS License Management Utility Manual*
- If DECnet for OpenVMS software is not included in your system, ignore the message; it is normal and does not affect the UETP run.

If you encounter other DECnet related errors, you should do the following:

- Run DECnet for OpenVMS software as a single phase (see Section 17.4.1) to determine whether the error can be re-created.
- Use the Help Message or refer to the *OpenVMS System Messages and Recovery Procedures Reference Manual*.

Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

17.7.10 Errors Logged but Not Displayed

If no errors are displayed at the console terminal or reported in the UETP.LOG file, you should run ERROR LOG to see if any errors were logged in the ERRLOG.SYS file. See the *OpenVMS System Management Utilities Reference Manual* for information about running the ERROR LOG.

17.7.11 No PCB or Swap Slots

The following error message indicates that no PCB or swap slots are available:

```
%UETP-I-BEGIN, UETLOAD00 beginning at 22-JUN-1995 07:47:16.50
%UETP-I-BEGIN, UETLOAD02_0000 beginning at 22-JUN-1995 07:47:16.76
%UETP-I-BEGIN, UETLOAD03_0001 beginning at 22-JUN-1995 07:47:16.92
%UETP-I-BEGIN, UETLOAD04_0002 beginning at 22-JUN-1995 07:47:17.13
%UETP-I-BEGIN, UETLOAD05_0003 beginning at 22-JUN-1995 07:47:17.35
%UETP-I-BEGIN, UETLOAD06_0004 beginning at 22-JUN-1995 07:47:17.61
%UETP-W-TEXT, The process -UETLOAD07_0005- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD08_0006- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD09_0007- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD10_0008- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD11_0009- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-ABORT, UETLOAD00 aborted at 22-JUN-1995 07:47:54.10
-UETP-W-TEXT, Aborted via a user Ctrl/C.
*****
*
*   END OF UETP PASS 1 AT 22-JUN-1995 07:48:03.17
*
*****
```

Solution

To solve this problem, use the following procedure:

1. Individually rerun the phase that caused the error message (the LOAD phase in the previous example) to see if the error can be reproduced.
2. Increase the size of the page file, using either the command procedure SYS\$UPDATE:SWAPFILES.COM (see Chapter 15) or SYSGEN (see the *OpenVMS System Management Utilities Reference Manual*).
3. Increase the system parameter MAXPROCESSCNT, if necessary.
4. Reboot the system.

17.7.12 No Keyboard Response or System Disk Activity

If the keyboard does not respond or the system disk is inactive, the system might be hung.

Solution

A system hangup can be difficult to trace; you should save the dump file for reference. To learn why the system hung, run the System Dump Analyzer as described in the *OpenVMS VAX System Dump Analyzer Utility Manual*.

Testing the System with UETP

17.7 Troubleshooting: Possible UETP Errors

Reasons for a system hangup include the following:

- Insufficient pool space—Increase the value of the system parameter NPAGEVIR and reboot the system.
- Insufficient page file space—Increase the page file space using the SYSGEN as described in the *OpenVMS System Management Utilities Reference Manual*.
- I/O device failure causing driver-permanent loop—Call Digital Services.

17.7.13 Lack of Default Access for the FAL Object

If default FAL access is disabled at the remote node selected by UETP for DECnet testing (the adjacent node on each active circuit, or a node defined by the group logical name UETP\$NODE_ADDRESS), messages similar to the following will appear:

```
%UETP-W-TEXT, The process -SVA019841_0001- returned a final status of:
%COPY-E-OPENOUT, error opening !AS as output
```

These messages are followed by:

```
%COPY-E-OPENOUT, error opening 9999""::SVA019841.D1; as output
-RMS-E-CRE, ACP file create failed
-SYSTEM-F-INVLOGIN, login information invalid at remote node
%COPY-W-NOTCOPIED, SYSSCOMMON:[SYSTEST]UETP.COM;2 not copied
%UETP-E-TEXT, Remote file test data error
```

You can ignore these messages.

17.7.14 Bugchecks and Machine Checks

When the system aborts its run, a bugcheck message appears at the console.

Solution

Call Digital Services. Often a hardware problem causes bugchecks and machine checks; solving bugchecks or machine checks is not easy. However, saving the SYSS\$SYSTEM:SYSDUMP.DMP and ERRLOG.SYS files is important so they are available for examination. Knowing whether the failure can be re-created is also important; you can run UETP again to verify the failure.

17.8 UETP Tests and Phases

This section explains, in detail, the organization of UETP and the individual components within the test package. You run UETP by starting a master command procedure containing commands to start each test phase. The procedure begins by prompting you for information needed by the various test phases. (See Section 17.4 for a detailed description of starting UETP.)

The master command procedure, UETP.COM, contains commands that initiate each test phase. UETP.COM also contains commands that do such tasks as defining logical names and manipulating files generated by the tests.

The UETP.COM procedure also issues commands to start the test controlling program UETPHAS00.EXE, which, in turn, controls each test phase. The test controller starts up multiple detached processes. It also reports their completion status and other information the processes report to it.

The sections that follow describe the various UETP test phases.

17.8.1 Initialization Phase

The following occurs during the initialization phase:

- The image UETINIT00.EXE prompts you for information. (See Section 17.4.) Your information defines variables that affect the execution of UETP tests.
- The image UETINIT01.EXE gathers information on all the controllers in the system and on their associated devices. This image writes the information into a file called UETINIDEV.DAT.
- Using the information in UETSUPDEV.DAT, UETINIT01.EXE verifies which devices in the configuration are operable by running the appropriate device test. Each device test completes a simple read/write operation to each device. If a device fails this test, the device's entry in UETINIDEV.DAT specifies that the device cannot be tested. As a result, subsequent UETP tests ignore that device.
- For each testable controller, UETINIT01.EXE writes a line into a file called UETCONT00.DAT. The line associates a test file with the controller it tests.

A summary of UETINIDEV.DAT always exists in UETP.LOG, and UETINIT01.EXE sends that summary to the console if you have requested the long report format.

17.8.2 Device Test Phase

The device test phase includes separate tests for each type of device, such as disk, magnetic tape, line printer, and terminal. This section explains the device test phase and presents instructions for testing a single device. If you want to run the entire device test phase individually, refer to Section 17.4.1.

17.8.2.1 How the Device Phase Works

The UETP device test phase starts an executable image, the phase controller UETPHAS00, which creates a detached process for every device controller to be tested. For example, if a system includes three terminal controllers, one line printer controller, and two disk controllers, the image creates six detached processes. In parallel, the detached processes execute images that test the various types of devices.

The initialization phase of UETP creates a file called UETINIDEV.DAT and a file called UETCONT00.DAT. UETINIDEV.DAT contains data on the controllers in the system supported by OpenVMS and their associated devices; UETCONT00.DAT associates a device test image with each testable controller.

UETPHAS00 uses the information in UETCONT00.DAT to find a device controller name to pass to each detached process that it creates. UETPHAS00 passes the controller name by writing it to a mailbox that is SYS\$INPUT to individual tests. Each detached process uses that data to determine which controller to test. The test image then searches UETINIDEV.DAT for the device controller and for all testable units on that controller. The phase controller terminates when all devices on all controllers have completed testing.

Because UETCONT00.DAT is deleted automatically at the end of a UETP run, you cannot run the device phase unless you start UETP.COM; you can run only individual test images. UETINIDEV.DAT exists in SYS\$TEST unless you delete it.

Testing the System with UETP

17.8 UETP Tests and Phases

17.8.2.2 Running a Single Device Test

You must be logged in to the SYSTEST account to run the individual tests as described in this section. Also, a copy of UETINIDEV.DAT must exist. If a copy of the file is not present from a previous run (a run of the entire UETP or a run of the device test phase creates UETINIDEV.DAT), you can create it. Note that when you run a single test, no log file is created; the test sends all its output to your terminal.

If you do not want to test all the device types, you can test a specific controller by choosing a test image name from Table 17-1 (for VAX systems) or Table 17-2 (for Alpha systems) and executing it as in the following example:

```
$ RUN UETTTYS00
Controller designation?: TTB
```

UETP prompts you for the controller designation and the device code. Unless you are testing your own terminal, you must explicitly designate a controller name. If you are running the terminal test, you can press Return to test your terminal only.

If you plan to repeat the run several times, you might find it more convenient to define the logical name CTRLNAME as follows:

```
$ DEFINE CTRLNAME TTB
$ RUN UETTTYS00
```

When you define the controller name in this way, the logical name CTRLNAME remains assigned after the test completes. To deassign this logical name, use the DCL command DEASSIGN as follows:

```
$ DEASSIGN CTRLNAME
```

17.8.2.3 Format of UETINIDEV.DAT

The UETINIDEV.DAT file is an ASCII sequential file that you can type or edit if necessary. The contents of this file are shown in the following command sequence:

```
$ TYPE UETINIDEV.DAT

DDB x ddd
UCB y uuuuu nnnnnnnnn.nnn
END OF UETINIDEV.DAT
```

The symbols in this example are defined as follows:

Symbol	Value
x	T, if testable units exist for this controller; N, if this controller is not to be tested
y	T, if this unit is testable; N, if this unit is not testable
ddd	Device controller name, for example DUA
uuuuu	Device unit number, for example 25
nnnnnnnnn.nnn	UETP device test name for the unit, for example, UETDISK00.EXE

UETINIDEV.DAT contains a DDB (device data block) line for each controller connected or visible to your system. After the DDB line is a UCB (unit control block) line for each unit connected to that controller. A device test can test a particular device only if both the DDB line and the UCB line indicate that the device is testable.

17.8.2.4 Running a Test in Loop Mode

If you want to put extra stress on a device, you can run the device test in loop mode, which causes the test to run indefinitely. For example:

```
$ DEFINE MODE LOOP
$ RUN UETDISK00
Controller designation?: DRA
%UETP-I-TEXT, End of pass 1 with 980 iterations at 22-JUN-1995 16:18:51:03
^C
```

You must use Ctrl/C to terminate the test run. If you use Ctrl/Y, UETP does not complete cleanup procedures.

17.8.2.5 Functions of Individual Device Tests

For each disk in the system, the disk test allocates two files into which it randomly writes blocks of data. The test then checks the data, reports any errors to SYS\$OUTPUT, and deletes the disk files.

When you run the disk test phase in a cluster environment, the test accesses all disks that are mounted by the system being tested, and users of the disk being tested might encounter an insufficient disk space problem. You should warn users on remote nodes (who share disks with users on the local system) that UETP might be testing a disk they are using.

The magnetic tape test exercises all the magnetic tape drives in the system. The test creates a large file on each mounted magnetic tape, into which it writes multiple sequential records of varying sizes. After writing the records, the test rewinds the magnetic tape, validates the written records, and reinitializes the magnetic tape.

The terminal and line printer test generates several pages or screens of output, in which each page or screen contains a header line and a test pattern of ASCII characters. A header line contains the test name, the device name, the date, and the time.

VAX

For the laboratory peripheral accelerator (LPA11-K), the test image determines the configuration on the LPA11-K's I/O bus. The image loads all types of microcode to the LPA11-K and reads or writes data for each device on the LPA11-K I/O bus. ♦

The communications device tests fill the transmit message buffer with random data; then, using loopback mode, the tests transmit and receive the message several times. To check that the looped-back data is correct, an AST routine is associated with a \$QIO read to compare the received message against the transmitted message. The procedure is repeated using messages of different lengths.

The interface device tests put the devices they are testing in maintenance mode, write random data, and then verify the data.

The Ethernet adapter test does self-test diagnostics on the device. It also does read and write tasks with test data that uses various adapter modes (such as internal loopback and external loopback).

The vector processor device test performs simple vector-scalar and vector-vector arithmetic operations and compares the results with expected values. The test also uses vector-related system service extensions and forces the system to generate arithmetic and memory management exceptions.

Testing the System with UETP

17.8 UETP Tests and Phases

VAX

Table 17-1 lists the device test images and the devices to be tested on VAX systems.

Table 17-1 Device Tests (VAX Only)

Test Image Name	Devices Tested
UETDISK00.EXE	Disks
UETTAPE00.EXE	Magnetic tape drives and tape cartridge drives
UETTTY00.EXE	Terminals and line printers
UETLPAK00.EXE	LPA11-K
UETCOMS00.EXE	DMC11, DMR11
UETDMPF00.EXE	DMF32, DMP11
UETDR1W00.EXE	DR11-W
UETDR7800.EXE	DR780, DR750
UETCDRO00.EXE	RRD40, RRD42, RRD50
UETUNAS00.EXE	Ethernet Adapters
UETVECTOR.EXE	Vector Processor, VVIEF

Alpha

Table 17-2 lists the device test images and the devices to be tested on Alpha systems.

Table 17-2 Device Tests (Alpha Only)

Test Image Name	Devices Tested
UETDISK00.EXE	Disks
UETTAPE00.EXE	Magnetic tape drives and tape cartridge drives
UETTTY00.EXE	Terminals and line printers
UETCDRO00.EXE	RRD42
UETUNAS00.EXE	Ethernet adapters

17.8.3 System Load Test Phase

The purpose of the system load test is to simulate a number of terminal users who are demanding system resources simultaneously. The system load tests, directed by the file UETLOAD00.DAT, create a number of detached processes that execute various command procedures. Each process simulates a user logged in at a terminal; the commands within each procedure are the same types of commands that a user enters from a terminal. The load test creates the detached processes in quick succession, and the processes generally execute their command procedures simultaneously. The effect on the system is analogous to an equal number of users concurrently issuing commands from terminals. In this way, the load test creates an environment that is similar to normal system use.

The load test uses the logical name LOADS to determine the number of detached processes to create. When you initiate the UETP command procedure, it prompts for the number of users to be simulated (see Section 17.4.3) and consequently the number of detached processes to be created. Your response, which depends on the amount of memory and the swapping and paging space in your system, defines the group logical name LOADS.

The UETP master command procedure deassigns all group logical names assigned by its tests as part of the termination phase. The group logical name LOADS remains assigned only if the UETP package does not complete normally.

The command procedures executed by the load test can generate a large amount of output, depending on the number of detached processes created. For each detached process (or user), the test creates a version of an output file called UETLOnnnn.LOG (*nnnn* represents a string of numeric characters). The console displays only status information as the load test progresses.

Whether the load test runs as part of the entire UETP or as an individual phase, UETP combines the UETLOnnnn.LOG files, writes the output to the file UETP.LOG, and deletes the individual output files.

You can run the system load test as a single phase by selecting LOAD from the choices offered in the startup dialog. (See Section 17.4.1.)

17.8.4 DECnet for OpenVMS Test Phase

If DECnet for OpenVMS software is included in your OpenVMS system, a run of the entire UETP automatically tests DECnet hardware and software. Because communications devices are allocated to DECnet and the DECnet devices cannot be tested by the UETP device test, UETP will not test the Ethernet adapter if DECnet for OpenVMS or another application has allocated the device. The DECnet node and circuit counters are zeroed at the beginning of the DECnet test to allow for failure monitoring during the run.

As with other UETP phases, you can run the DECnet for OpenVMS phase individually by following the procedure described in Section 17.4.1.

17.8.4.1 Environment

The DECnet for OpenVMS test will work successfully on OpenVMS systems connected to all DECnet supported node types, including routing and nonrouting nodes and several different types of operating systems (such as RSTS, RSX, TOPS, and RT). To copy files between systems, the remote systems must have some type of default access. The DECnet phase tests the following:

- The node on which UETP is running.
- All circuits in sequence, unless you have defined the logical name UETP\$NODE_ADDRESS to be the remote node that you want to run the test on. If you have defined a remote node, the DECnet phase tests only one circuit.
- All adjacent or first-hop nodes and all circuits in parallel.

No limit exists on the number of communication lines supported by the tests. A test on one adjacent node should last no more than two minutes at normal communications transfer rates.

Note

UETP assumes your system has default access for the FAL object, even though the network configuration command procedure NETCONFIG.COM does not provide access for the FAL object by default. When you install DECnet software with the defaults presented by NETCONFIG.COM, the UETP DECnet phase can produce error messages. You can ignore these error messages. See Section 17.7.13 for more information.

Testing the System with UETP

17.8 UETP Tests and Phases

17.8.4.2 How the DECnet Phase Works

UETP (under the control of UETPHAS00.EXE) reads the file UETDNET00.DAT and completes the following steps during the DECnet for OpenVMS phase:

1. Executes a set of Network Control Program (NCP) LOOP EXECUTOR commands to test the node on which UETP is running.
2. Uses NCP to execute the command SHOW ACTIVE CIRCUITS. The results are placed in UETININET.TMP, from which UETP creates the data file UETININET.DAT. The UETININET.TMP file contains the following information for any circuit in the ON state but not in transition:

- Circuit name
- Node address
- Node name (if one exists)

The UETININET.TMP file is used throughout the DECnet phase to determine which devices to test.

3. Uses the UETININET.TMP file to create an NCP command procedure for each testable circuit. Each command procedure contains a set of NCP commands to zero the circuit and node counters and to test the circuit and adjacent node by copying files back and forth.

Note

If you do not want the counters zeroed, do not test the DECnet for OpenVMS software.

4. Executes the command procedures from Step 3 in parallel to simulate a heavy user load. The simulated user load is the lesser of the following values:
 - The number of testable circuits, multiplied by two
 - The maximum number of user-detached processes that can be created on the system before it runs out of resources (determined by UETINIT00)
5. Executes a program, UETNETS00.EXE, that uses the UETININET.DAT file to check the circuit and node counters for each testable circuit. If a counter indicates possible degradation (by being nonzero), its name and value are reported to the console. All counters are reported in the log file, but only the counters that indicate degradation are reported to the console. Following is an example of UETNETS00 output:

```
%UETP-S-BEGIN, UETNETS00 beginning at 22-JUN-1995 13:45:33.18
%UETP-W-TEXT, Circuit DMC-0 to (NODENAME1) OK.
%UETP-I-TEXT, Node (NODENAME2) over DMC-1 response timeouts = 1.
%UETP-I-TEXT, Circuit DMC-1 to (NODENAME2) local buffer errors = 34.
%UETP-I-TEXT, Node (NODENAME3) over DMP-0 response timeouts = 3.
%UETP-S-ENDED, UETNETS00 ended at 22-JUN-1995 13:45:36.34
```

Because degradation is not necessarily an error, the test's success is determined by you, not by the system. The following counters indicate possible degradation:

For Circuits

- Arriving congestion loss
- Corruption loss

- Transit congestion loss
- Line down
- Initialization failure
- Data errors inbound
- Data errors outbound
- Remote reply timeouts
- Local reply timeouts
- Remote buffer errors
- Local buffer errors
- Selection timeouts
- Remote process errors
- Local process errors
- Locally initiated resets
- Network initiated resets

For Nodes

- Response timeouts
- Received connect resource errors
- Node unreachable packet loss
- Node out of range packet loss
- Oversized packet loss
- Packet format error
- Partial routing update loss
- Verification reject

17.8.5 Cluster-Integration Test Phase

The cluster-integration test phase consists of a single program and a command file that depend heavily on DECnet for OpenVMS software. This phase uses DECnet for OpenVMS software to create SYSTEST_CLIG processes on each OpenVMS node in the cluster and to communicate with each node. SYSTEST_CLIG is an account that is parallel to SYSTEST, but limited so that it can only be used as part of the cluster-integration test. The following restrictions on the SYSTEST_CLIG account are necessary for a correct run of the cluster test phase:

- The account must be enabled and the password must be null. For more information, see Section 17.3.16.
- The UIC must be the same as that of the SYSTEST account.
- The account must have the same privileges and quotas as the SYSTEST account. For more information, see Section 17.7.2.
- The account can allow login only through DECnet for OpenVMS software.
- The account must be locked into running UETCLIG00.COM when it logs in.

Testing the System with UETP

17.8 UETP Tests and Phases

These items are necessary to ensure the security and privacy of your system. If the test cannot create a SYSTEST_CLIG process on an OpenVMS node, it gives the reason for the failure and ignores that node for the lock tests and for sharing access during the file test. Also, the test does not copy log files from any node on which it cannot create the SYSTEST_CLIG process. If a communication problem occurs with a SYSTEST_CLIG process after the process has been created, the test excludes the process from further lock and file sharing tests. At the end of the cluster-integration test, an attempt is made to report any errors seen by that node.

UETCLIG00.EXE has two threads of execution: the primary and the secondary. The first, or primary thread, checks the cluster configuration (OpenVMS nodes, HSC nodes, and the attached disks that are available to the node running the test). For selected OpenVMS nodes, the primary thread attempts to start up a SYSTEST_CLIG process through DECnet software. If the primary thread was able to start a SYSTEST_CLIG process on a node, the node runs the command file UETCLIG00.COM, which starts up UETCLIG00.EXE and runs the secondary execution thread.

The process running the primary thread checks to see that it can communicate with the processes running the secondary threads. It then instructs them to take out locks so that a deadlock situation is created.

The primary thread tries to create a file on some disk on selected OpenVMS and HSC nodes in the cluster. It writes a block, reads it back, and verifies it. Next, it selects one OpenVMS node at random and asks that node to read the block and verify it. The primary thread then extends the file by writing another block and has the secondary thread read and verify the second block. The file is deleted.

The secondary processes exit. They copy the contents of their SYS\$ERROR files to the primary process, so that the UETP log file and console report show all problems in a central place. DECnet for OpenVMS software automatically creates a NETSERVER.LOG in SYS\$TEST as the test is run, so that if necessary, you can read that file later from the node in question.

During the test run, the primary process uses the system service SYS\$BRKTHRU to announce the beginning and ending of the test to each OpenVMS node's console terminal.

You can define the group logical name MODE to the equivalence string DUMP to trace most events as they occur. Note that the logical name definitions apply only to the node on which they were defined. You must define MODE on each node in the VMScluster on which you want to trace events.

Getting Information About the System

This chapter discusses setting up and maintaining system log files, maintaining error log files, and using system management utilities to monitor the system.

This chapter describes the following tasks:

Task	Section
Using the Error Formatter (ERRFMT)	Section 18.3
Using the Error Log utility (ERROR LOG) to produce reports	Section 18.4
‡Using DECEvent to report system events	Section 18.5
Setting up, maintaining, and printing the operator log file	Section 18.6
Using security auditing	Section 18.7
Using the Monitor utility to monitor system performance	Section 18.8
‡Alpha specific	

This chapter explains the following concepts:

Concept	Section
System log files	Section 18.1
Error logging	Section 18.2
Error Log utility (ERROR LOG)	Section 18.4.1
‡DECEvent Event Management utility	Section 18.5.1
Operator log file	Section 18.6.1
OPCOM messages	Section 18.6.2
Security auditing	Section 18.7.1
Monitor utility (MONITOR)	Section 18.8.1
‡Alpha specific	

18.1 Understanding System Log Files

In maintaining your system, collect and review information about system events. The operating system provides several log files that record information about the use of system resources, error conditions, and other system events. Table 18–1 briefly describes each file and provides references to sections that discuss the files in more detail.

Getting Information About the System

18.1 Understanding System Log Files

Table 18-1 System Log Files

Log File	Description	For More Information
Error log file	The system automatically records device and CPU error messages in this file.	See Section 18.2.
Operator log file	The operator communication manager (OPCOM) records system events in this file.	See Chapter 2, Section 18.6, and Section 19.6.
Accounting file	The accounting file tracks the use of system resources.	See Chapter 19.
Security audit log file	The audit server process preallocates disk space to and writes security-relevant system events to this file.	See Section 18.7.

18.2 Understanding Error Logging

The error logging facility automatically writes error messages to the latest version of the error log file, SYS\$ERRORLOG:ERRLOG.SYS. You can use the Error Log utility (ERROR LOG) to report selectively on error log files.

Alpha

On Alpha systems, you must use the DECEvent Event Management utility to produce reports derived from system event entries. ♦

Error log reports are primarily intended for use by Digital Services personnel to identify hardware problems. System managers often find error log reports useful in identifying recurrent system failures that require outside attention.

Parts of the Error Logging Facility

The error logging facility consists of the parts shown in Table 18-2.

Table 18-2 Parts of the Error Logging Facility

Part	Description
Executive routines	Detect errors and events, and write relevant information into error log buffers in memory.
Error Formatter (ERRFMT)	The ERRFMT process, which starts when the system is booted, periodically empties error log buffers, transforms the descriptions of errors into standard formats, and stores formatted information in an error log file on the system disk. (See Section 18.3.2.) The Error Formatter allows you to send mail to the SYSTEM account or another user if the ERRFMT process encounters a fatal error and deletes itself. (See Section 18.3.3.)
Error Log utility (ERROR LOG)	Invokes the Error Log Report Formatter (ERF) , which selectively reports the contents of an error log file. You invoke ERROR LOG by entering the DCL command ANALYZE/ERROR_LOG. (See Section 18.4.2.)
‡DECEvent	Selectively reports the contents of an event log file; you invoke DECEvent by entering the DCL command DIAGNOSE. (See Section 18.5.)
‡Alpha specific	

The executive routines and the Error Formatter (ERRFMT) process operate continuously without user intervention. The routines fill the error log buffers in memory with raw data on every detected error and event. When one of the

available buffers becomes full, or when a time allotment expires, ERRFMT automatically writes the buffers to SYS\$ERRORLOG:ERRLOG.SYS.

Sometimes a burst of errors can cause the buffer to fill up before ERRFMT can empty them. You can detect this condition by noting a skip in the error sequence number of the records reported in the error log reports. As soon as ERRFMT frees the buffer space, the executive routines resume preserving error information in the buffers.

The ERRFMT process displays an error message on the system console terminal and stops itself if it encounters excessive errors while writing the error log file. Section 18.3.1 explains how to restart the ERRFMT process.

18.3 Using the Error Formatter (ERRFMT)

The ERRFMT process is started automatically at boot time. The following sections explain how to perform these tasks:

Task	Section
Restart the ERRFMT process, if necessary	Section 18.3.1
Maintain error log files	Section 18.3.2
Send mail if the ERRFMT process is deleted	Section 18.3.3

18.3.1 Restarting the ERRFMT Process

To restart the ERRFMT process, follow these steps:

1. Log in to the system manager's account so that you have the required privileges to perform the operation.
2. Execute the site-independent startup command procedure (STARTUP.COM), specifying ERRFMT as the command parameter, as follows:

```
$ @SYS$SYSTEM:STARTUP ERRFMT
```

Note

If disk quotas are enabled on the system disk, ERRFMT starts only if UIC [1,4] has sufficient quotas.

18.3.2 Maintaining Error Log Files

Because the error log file, SYS\$ERRORLOG:ERRLOG.SYS, is a shared file, ERRFMT can write new error log entries while the Error Log utility reads and reports on other entries in the same file.

ERRLOG.SYS increases in size and remains on the system disk until you explicitly rename or delete it. Therefore, devise a plan for regular maintenance of the error log file. One method is to rename ERRLOG.SYS on a daily basis. If you do this, the system creates a new error log file. You might, for example, rename the current copy of ERRLOG.SYS to ERRLOG.OLD every morning at 9:00. To free space on the system disk, you can then back up the renamed version of the error log file on a different volume and delete the file from the system disk.

Getting Information About the System

18.3 Using the Error Formatter (ERRFMT)

Another method is to keep the error log file on a disk other than the system disk by defining the logical name SYS\$ERRORLOG to be the device and directory where you want to keep error log files; for example:

```
$ DEFINE/SYSTEM/EXECUTIVE SYS$ERRORLOG DUA2:[ERRORLOG]
```

To define this logical name each time you start up the system, add the logical name definition to your SYLOGICALS.COM procedure. See Section 5.2.5 for details.

Be careful not to delete error log files inadvertently. You might also want to adopt a file-naming convention that includes a beginning or ending date for the data in the file name.

18.3.3 Using ERRFMT to Send Mail

The Error Formatter (ERRFMT) allows you to send mail to the system manager or to another designated user if the ERRFMT process encounters a fatal error and deletes itself.

Two system logical names, ERRFMT\$_SEND_MAIL and ERRFMT\$_SEND_TO, control this feature:

- **ERRFMT\$_SEND_MAIL**

To enable sending mail, must translate to the string TRUE, and is case insensitive. Any other value disables the sending of mail.

- **ERRFMT\$_SEND_TO**

Must translate to a user name (the current default is SYSTEM).

Digital recommends that you do not use distribution lists and multiple user names.

You can define these logical names in one of two ways:

- Dynamically, using DCL DEFINE/SYSTEM commands

After you make the changes, you must stop and restart ERRFMT for the changes to take effect.

- Permanently, in SYS\$STARTUP:SYLOGICAL.COM

The logical names you define take effect the next time the system is rebooted. The following instructions use this method.

18.3.3.1 Enabling and Disabling ERRFMT to Send Mail

If ERRFMT\$_SEND_MAIL is defined to be TRUE, the system manager receives a mail message that contains a subject line saying that ERRFMT is about to delete itself. The operator log file and the output displayed at the system console, OPA0;, contain more detailed information about the failure encountered and instructions on how to restart ERRFMT; however, system managers are often not at the console to see this information.

If you are using ERRFMT in one mode, for example, with sending mail enabled, and you want to disable sending mail, use the system manager's account to edit SYS\$STARTUP:SYLOGICAL.COM, adding the following command:

```
$ DEFINE/SYSTEM ERRFMT$_SEND_MAIL FALSE
```

To reenable sending mail, use the system manager's account to edit SYS\$STARTUP:SYLOGICAL.COM, adding the following command:

```
$ DEFINE/SYSTEM ERRFMT$_SEND_MAIL TRUE
```


Getting Information About the System

18.3 Using the Error Formatter (ERRFMT)

18.3.3.2 Sending Mail to Another User

Sending mail to the SYSTEM account is enabled by default. However, you can define `ERRFMT$_SEND_TO` to send mail to another user if `ERRFMT` is about to delete itself.

To change the user name to receive mail, use the system manager's account to edit `SYS$STARTUP:SYLOGICAL.COM`, adding an appropriate logical name `DEFINE` command. For example:

```
$ DEFINE/SYSTEM ERRFMT$_SEND_TO R_SMITH
```

Digital recommends that you do not use distribution lists and multiple user names.

18.4 Using the Error Log Utility (ERROR LOG)

You use the Error Log utility (`ERROR LOG`) to report selectively on the contents of an error log file. You must have `SYSPRV` to run `ERROR LOG`.

18.4.1 Understanding the Error Log Utility (ERROR LOG)

`ERROR LOG` supports most OpenVMS-supported hardware, such as adapters, disks, tapes, CPUs, and memories, but not all communications devices. Some synchronous communications devices are supported.

The operating system automatically writes messages to the latest version of an error log file, `SYS$ERRORLOG:ERRLOG.SYS`, as the events shown in Table 18–3 occur.

Table 18–3 Types of Events Reported in the Error Log File

Event	Description
Errors	Device errors, device timeouts, machine checks, bus errors, memory errors (hard or soft error correcting code [ECC] errors), asynchronous write errors, and undefined interrupts
Volume changes	Volume mounts and dismounts
System events	System startups, messages from the Send Message to Error Logger (<code>\$SENDERR</code>) system service, and time stamps

You can use `ERROR LOG` to process error log entries for the following forms of optional output:

- Full report of selected entries, which is the default
- Brief report of selected entries
- Summary report of selected entries
- Register dump report of selected device entries
- Binary copy of selected entries
- Binary copy of rejected entries

Section 18.4.2 explains how to produce error log reports. See the *OpenVMS System Management Utilities Reference Manual* for examples of error log reports.

The error reports that `ERROR LOG` produces are useful in two ways:

- They aid preventive maintenance by identifying areas within the system that show potential for failure.

Getting Information About the System

18.4 Using the Error Log Utility (ERROR LOG)

- They aid the diagnosis of a failure by documenting the errors and events that led up to it.

The detailed contents of the reports are most meaningful to Digital Services personnel. However, you can use the reports as an important indicator of the system's reliability. For example, using the DCL command `SHOW ERROR`, you might see that a particular device is producing a relatively high number of errors. You can then use `ERROR LOG` to obtain a more detailed report and decide whether to consult Digital Services. If you do, Digital Services personnel can run diagnostic programs to investigate the device and attempt to isolate the source of the errors.

If a system component fails, a Digital Services representative can study the error reports of the system activity leading up to and including the failure. If a device fails, you can generate error reports immediately after the failure; for example:

- One report might describe in detail all errors associated with the device that occurred within the last 24 hours.
- Another report might summarize all types of errors for all devices that occurred within the same time period.
- The summary report can put the device errors into a systemwide context.

The Digital Services representative can then run the appropriate diagnostic program for a thorough analysis of the failed device. Using the combined error logging and diagnostic information, the Digital Services representative can proceed to correct the device.

Error reports allow you to anticipate potential failures. Effective use of the Error Log utility in conjunction with diagnostic programs can significantly reduce the amount of system downtime.

18.4.2 Producing Error Log Reports

You enter the DCL command in the following format:

```
ANALYZE/ERROR_LOG [/qualifier(s)][file_spec,...]
```

where:

<i>qualifier</i>	Specifies the function the <code>ANALYZE/ERROR_LOG</code> command is to perform.
<i>file-spec</i>	Specifies one or more files that contain information to be interpreted for the error log report.

See the *OpenVMS System Management Utilities Reference Manual* for details about the command and its parameters and for examples of error log reports.

`ERROR LOG` issues error messages for inconsistent error log entries. The *OpenVMS System Messages and Recovery Procedures Reference Manual* lists these messages and provides explanations and suggested user actions.

18.4.3 Producing a Full Error Log Report

The following steps show how to produce an error log report for all entries in the error log file and how to print the report:

1. Either log in to the `SYSTEM` account or ensure that you have the `SYSPRV` privilege. (You must have privilege to access the error log file.) For example:

```
$ SET PROCESS/PRIVILEGE=SYSPRV
```


Getting Information About the System

18.4 Using the Error Log Utility (ERROR LOG)

2. Set your default disk and directory to SYS\$ERRORLOG:
\$ SET DEFAULT SYS\$ERRORLOG
3. Examine the error log directory to see which error log file you want to analyze:
\$ DIRECTORY
4. To obtain a full report of the current error log file, enter the following command:
\$ ANALYZE/ERROR_LOG/OUTPUT=ERRORS.LIS
5. Print a copy of the report, using the file name you specified with the /OUTPUT qualifier:
\$ PRINT ERRORS.LIS

Example

```
$ SET PROCESS/PRIVILEGE=SYSPRV
$ SET DEFAULT SYS$ERRORLOG
❶ $ DIRECTORY
Directory SYS$SYSROOT:[SYSERR]
ERRLOG.OLD;2 ERRLOG.OLD;1 ERRLOG.SYS;1
Total of 3 files.
❷ $ ANALYZE/ERROR_LOG/OUTPUT=ERRORS.LIS ERRLOG.OLD
❸ $ PRINT ERRORS.LIS
```

Following are explanations of the commands in the example.

- ❶ The DIRECTORY command lists all the files in the SYS\$ERRORLOG directory. The directory contains three files: two old error log files and the current error log file, ERRLOG.SYS.
- ❷ The ANALYZE/ERROR_LOG command writes a full report to a file called ERRORS.LIS, using the most recent ERRLOG.OLD file as input.
- ❸ The PRINT command prints ERRORS.LIS.

18.4.4 Using Other Error Log Report Options

This section briefly explains how to specify report formats and produce a report of selected entries.

Table 18-4 contains error log report options. For more details about options and examples of error log reports using options, see the *OpenVMS System Management Utilities Reference Manual*.

Getting Information About the System

18.4 Using the Error Log Utility (ERROR LOG)

Table 18–4 Error Log Report Options

In Order To...	You Can...
Specify report formats	<p>Change report formats by using qualifiers, including the following:</p> <ul style="list-style-type: none">• <code>/BINARY</code>—to convert binary error log records to ASCII text or to copy error log records to a specified output file.• <code>/BRIEF</code>—to create a brief report.• <code>/REGISTER_DUMP</code>—to generate, in a hexadecimal longword format, a report that consists of device register information (used in conjunction with the <code>/INCLUDE</code> qualifier).• <code>/REJECTED</code>—to specify the name of a file that will contain binary records for rejected entries.
Specify a display device for reports	<p>Use the <code>/OUTPUT</code> qualifier to send reports to a terminal for display or to a disk or magnetic tape file. By default, the system sends the report to the <code>SYS\$OUTPUT</code> device. Because ERROR LOG reports are 72 columns wide, you can display them on the terminal screen.</p>
Produce a report of selected entries	<p>Use qualifiers to produce error log reports for specific types of events and for a specified time interval. For example, you can process error log entries by selecting a time interval using the <code>/SINCE</code>, <code>/BEFORE</code>, or <code>/ENTRY</code> qualifiers.</p> <p>You can specify error log entries for specific events by using the qualifiers <code>/INCLUDE</code> and <code>/EXCLUDE</code>. These qualifiers form a filter to determine which error log entries are selected or rejected.</p> <p>In addition, you can generate error log reports for one or more VMScluster members by using the <code>/NODE</code> qualifier.</p>
Exclude unknown error log entries	<p>By default, when <code>ANALYZE/ERROR_LOG</code> encounters an unknown device, CPU, or error log entry, the utility produces the entry in hexadecimal longword format. Exclude these entries from the report by specifying <code>/EXCLUDE=UNKNOWN_ENTRIES</code> in the command line.</p>

18.5 Using the DECevent Event Management Utility (DECevent) (Alpha Only)

Alpha

On Alpha systems, the DECevent Event Management utility (DECevent) provides the interface between a system user and the operating system's event log files.

18.5.1 Understanding DECevent (Alpha Only)

DECevent allows system users to produce ASCII reports derived from system event entries. The format of the ASCII reports depends on the command entered on the command language interpreter (CLI) with a maximum character limit of 255 characters.

DECevent uses the error log file, `SYS$ERRORLOG:ERRLOG.SYS`, as the default input file, unless you specify another input file.

Event reports are useful for determining preventive maintenance by helping to identify areas within the system showing potential failure. Event reports also aid in the diagnosis of a failure by documenting events that led to the failure.

The contents of the event reports are most meaningful to Digital support personnel. However, you can use the event reports as an indicator of system reliability. For example, while using the DCL command `SHOW ERROR`, you might see that a particular device is producing a higher than normal number of

Getting Information About the System

18.5 Using the DECEvent Event Management Utility (DECEvent) (Alpha Only)

events. You can use DECEvent to obtain various detailed reports and determine if you need to contact Digital Services.

If a system component fails, Digital Services can use the event reports to create a history of events leading up to and including the failure.

Used in conjunction with diagnostic programs, event reports significantly reduce the amount of system down time.

DECEvent Report Types

DECEvent produces five types of reports:

Report Type	Description
Full (default)	Provides a translation of all available information for each entry in the event log.
Brief	Provides a translation of key information for each entry in the event log.
Terse	Provides binary event information and displays register values and other ASCII messages in a condensed format.
Summary	Provides a statistical summary of the event entries in the event log.
Fast Error (FSTERR)	Provides a quick, one-line per-entry report of your event log for a variety of disk devices.

These report types are mutually exclusive; in other words, you can select only one report type in a command.

Section 18.5.5 contains examples of types of reports. The *OpenVMS System Management Utilities Reference Manual* contains additional examples of the types of reports produced by DECEvent.

The following sections explain how to use DECEvent:

Task	Section
Invoking and exiting DECEvent	Section 18.5.2
Using DECEvent qualifiers	Section 18.5.3
Using additional DECEvent commands	Section 18.5.4
Producing DECEvent reports	Section 18.5.5

In addition, restrictions are listed in Section 18.5.6.

18.5.2 Invoking and Exiting DECEvent (Alpha Only)

To invoke DECEvent, enter the following command:

```
$ DIAGNOSE/TRANSLATE [/qualifier(s)] [file-spec][,...]
```

Note

The /TRANSLATE qualifier is the default qualifier; typing it on the command line is not necessary.

DECEvent does not prompt you. To exit from DECEvent, press Ctrl/C and Return (otherwise, no prompt is returned).

Getting Information About the System

18.5 Using the DECevent Event Management Utility (DECevent) (Alpha Only)

18.5.3 Using DECevent Qualifiers (Alpha Only)

The DECevent qualifiers shown and described in Table 18–5 allow you to change the format of the reports that DECevent produces.

Table 18–5 DECevent Qualifiers (Alpha Only)

Qualifier	Description
/BEFORE	Specifies that only those entries dated earlier than the stated date and time are to be selected for the event report
/BINARY	Controls whether the binary error log records are converted to ASCII text or copied to the specified output file
/BRIEF	Generates a brief report
/CONTINUOUS	Specifies events are formatted in real time, as they are logged by the operating system event logger
/DUMP	Specifies the output to be a brief report followed by a dump of information from the input event log file
/ENTRY	Generates a report that includes the specified entry range or starts at the specified entry number
/EXCLUDE	Excludes events generated by the specified device class, device name, or error log entry type from the report
/FSTERR	Generates a quick, one-line-per-entry report for an event log entry for disks
/FULL	Generates a full report (default), which provides all available information for an event log entry
/INCLUDE	Includes events generated by the specified device class, device name, or error log entry type in the report
/INTERACTIVE	Allows users to exit from the command line interface and enter the DECevent interactive command shell
/LOG	Controls whether informational messages that specify the number of entries selected and rejected for each input file are sent to SYS\$OUTPUT
/NODE	Generates a report consisting of event entries for specific nodes in a VAXcluster system
/OUTPUT	Specifies the output file for the report
/REJECTED	Allows you to specify the name of a file that will contain binary records for rejected entries
/SINCE	Specifies that only those entries dated later than the stated date and time are to be selected for the report
/SUMMARY	Generates an event report that consists of a statistical summary
/TERSE	Generates an event report consisting of binary event information, register values and ASCII messages in a condensed format
/TRANSLATE	Is the default qualifier for the DIAGNOSE command verb

Do not use the /BINARY qualifier with any report type qualifier (/FULL, /BRIEF, /TERSE, /SUMMARY, and /FSTERR) or with the /OUTPUT qualifier.

18.5 Using the DECevent Event Management Utility (DECevent) (Alpha Only)

Privileges Required

- You must have SYSPRV privilege to run DECevent; however, only read access is required to access the ERRLOG.SYS file.
- You must have the DIAGNOSE privilege for the /CONTINUOUS qualifier to work, enabling the continuous display of events on a terminal screen.

18.5.4 Using Additional DECevent Commands (Alpha Only)

In addition to the qualifiers listed in Table 18–5, DECevent contains a set of DIRECTORY commands and a set of SHOW commands:

- DIRECTORY commands

These commands allow you to display a list of rulesets that DECevent needs to translate events into a readable format. (A **ruleset** is a software routine or function that is analogous to an executable file.)

The following DIRECTORY commands are currently implemented in DECevent:

- DIRECTORY EVENT

This command lists all rulesets associated with event translation.

- DIRECTORY CANONICAL

This command lists all rulesets associated with event reports.

- SHOW commands

These commands allow a user to view specific settings and selections. The following SHOW commands are currently implemented in DECevent:

- SHOW SELECT

By appending a specific selection keyword name to the SHOW SELECT command, you view only that selection keyword.

- SHOW SETTINGS

By appending a specific setting's name to the SHOW SETTINGS command, you view only that setting's name and value.

18.5.5 Producing DECevent Reports (Alpha Only)

This section contains examples of DECevent commands and reports.

18.5.5.1 Producing a Full Report (Alpha Only)

To produce a full report, use the /FULL qualifier. The full report format provides a translation of all available information for each entry in the event log. The full report is the default report type if a report type is not specified in the command line.

Both of the following commands will produce a full report format:

```
$ DIAGNOSE/TRANSLATE/FULL
```

```
$ DIAGNOSE
```

(/TRANSLATE and /FULL are defaults.)

Example 18–1 shows the format of a full report.

Getting Information About the System

18.5 Using the DECEvent Event Management Utility (DECEvent) (Alpha Only)

Example 18-1 Full Report Format (Alpha Only)

```
***** ENTRY 1 *****

Logging OS                      1. OpenVMS
System Architecture             2. Alpha
OS version                     V6.2
Event sequence number          1583.
Timestamp of occurrence         18-APR-1995 09:21:18
System uptime in seconds        58004.
Error mask                     x00000000
Flags                          x0001   Dynamic Device Recognition present
Host name                      COGENT

Alpha HW model                  DEC 3000 Model 400
System type register           x00000004 DEC 3000
Unique CPU ID                  x00000002
mpnum                          x000000FF
mperr                          x000000FF

Event validity                  -1. Unknown validity code
Event severity                  -1. Unknown severity code
Entry type                     100.
Major Event class              3. IO Subsystem

IO Minor Class                 1. MSCP
IO Minor Sub Class             5. Logged Message

---- Device Profile ----
Vendor
Product Name                   RAID 0 - Host Based
Unit Name                     COGENT$DPA
Unit Number                    10.
Device Class                   x0001   Disk

---- IO SW Profile ----
VMS DC$_CLASS                  1.
VMS DT$_TYPE                   175.

---- MSCP Logged Msg ----

Logged Message Type Code       22. RAID Message
RAID Event Type                8. Remove Member
Distinguished Member           0.
Member Index                   1.
RAID Urgency                   4. Global Disk Error
RAID Status                    x00180009 Bit 00 - Reduced
                                      Bit 03 - Striped
                                      Bit 19 - FE Dis FE
                                      Bit 20 - BC Buff Copy Off

RAIDset Name                   KGB
*****
```

18.5.2 Producing a Brief Report (Alpha Only)

To produce a brief report, use the /BRIEF qualifier. The brief report format provides translation of key information for each entry in the event log. For example:

```
$ DIAGNOSE/TRANSLATE/BRIEF
```


Getting Information About the System

18.5 Using the DECEvent Event Management Utility (DECEvent) (Alpha Only)

Example 18-2 shows the format of a brief report.

Example 18-2 Brief Report Format (Alpha Only)

```
***** ENTRY 1 *****

Logging OS                      1. OpenVMS
System Architecture              2. Alpha
OS version                      V6.2
Event sequence number           1583.
Timestamp of occurrence          18-APR-1995 09:21:18
System uptime in seconds         58004.
Error mask                      x00000000
Host name                       COGENT

Alpha HW model                  DEC 3000 Model 400
System type register            x00000004 DEC 3000
Unique CPU ID                   x00000002
mpnum                           x000000FF
mperr                           x000000FF

Event validity                  -1. Unknown validity code
Event severity                  -1. Unknown severity code
Major Event class               3. IO Subsystem

IO Minor Class                  1. MSCP
IO Minor Sub Class              5. Logged Message

---- Device Profile ----
Vendor
Product Name                    RAID 0 - Host Based
Unit Name                      COGENT$DPA
Unit Number                     10.
Device Class                    x0001 Disk

Logged Message Type Code        22. RAID Message
RAID Event Type                 8. Remove Member
Distinguished Member            0.
Member Index                    1.
RAID Urgency                    4. Global Disk Error
RAID Status                     x00180009 Bit 00 - Reduced
                                   Bit 03 - Striped
                                   Bit 19 - FE Dis FE
                                   Bit 20 - BC Buff Copy Off
RAIDset Name                    KGB
*****
```

18.5.5.3 Producing a Terse Report (Alpha Only)

To produce a terse report, use the /TERSE qualifier. The terse report format provides binary event information and displays register values and other ASCII messages in a condensed format. For example:

```
$ DIAGNOSE/TRANSLATE/TERSE
```

Example 18-3 shows the format of a terse report.

Getting Information About the System

18.5 Using the DECevent Event Management Utility (DECevent) (Alpha Only)

Example 18-3 Terse Report Format (Alpha Only)

```
***** ENTRY 1 *****
Logging OS 1.
System Architecture 2.
OS version V6.2
Event sequence number 1583.
Timestamp of occurrence 1995041809211800
System uptime in seconds 58004.
Error mask x00000000
Flags x0001
Host name COGENT
Alpha HW model DEC 3000 Model 400
System type register x00000004
Unique CPU ID x00000002
mpnum x000000FF
mperr x000000FF
Event validity -1.
Event severity -1.
Entry type 100.
Major Event class 3.
IO Minor Class 1.
IO Minor Sub Class 5.
---- Device Profile ----
Vendor
Product Name RAID 0 - Host Based
Unit Name COGENT$DPA
Unit Number 10.
Device Class x0001
---- IO SW Profile ----
VMS DC$_CLASS 1.
VMS DT$_TYPE 175.
---- MSCP Logged Msg ----
Logged Message Type Code 22.
RAID Event Type 8.
Distinguished Member 0.
Member Index 1.
RAID Urgency 4.
RAID Status x00180009
RAIDset Name KGB
*****
```

18.5.5.4 Producing a Summary Report (Alpha Only)

To produce a summary report, use the /SUMMARY qualifier. The summary report format provides a statistical summary of the event entries in the event log. For example:

```
$ DIAGNOSE/TRANSLATE/SUMMARY
```

Example 18-4 shows the format of a summary report.

Getting Information About the System

18.5 Using the DECevent Event Management Utility (DECevent) (Alpha Only)

Example 18-4 Summary Report Format (Alpha Only)

SUMMARY OF ALL ENTRIES LOGGED ON NODE COGENT

IO Subsystem

MSCP

9.

Host Based RAID

3.

DATE OF EARLIEST ENTRY

18-APR-1995 09:21:18

DATE OF LATEST ENTRY

12-MAY-1995 10:44:54

18.5.5.5 Producing a Fast Error (FSTERR) Report (Alpha Only)

To produce a Fast Error report, use the /FSTERR qualifier. For example:

```
$ DIAGNOSE/TRANSLATE/FSTERR
```

The Fast Error report provides a quick, one-line-per-entry report of your event log for a variety of disk devices. This makes event analysis and system troubleshooting much easier by eliminating extraneous event information. For example:

```
$ DIAGNOSE/FSTERR [infile]
```

A Fast Error report is shown in Example 18-5.

Example 18-5 Fast Error (FSTERR) Report Format (Alpha Only)

Drive Name	yymmdd	hhmmss	Entry	MSCP	Evnt	LED	LBN	Physical	HSC	Volume	Drive/
								Cyl Hd Sec RA	RP	Serial	
LUKE\$DUA070	921119	160754	3	00EB			255		70	71	V00717
LUKE\$DUA070	921119	160754	4	00EB			255		70	71	V00717
HSC015\$DUA028	910323	113204	5	00EB					70	51	V15039
HSC015\$DUA028	910323	113204	6	00EB					71	51	V15039
BATES\$DUA197	921118	002116	7	00EB					72	32	V17524
CHEWIE\$DUA101	911205	114908	8	00EB					73	81	V
PMASON\$DUA006	921207	165007	15	00EB			255		90	42	D23387
PMASON\$DUA006	921207	165007	16	00EB			255		90	42	D23387
C3P0\$DUA242	870218	060031	17	01AB					90	40	D48575
CHER\$DU2132*901008	231053		18	00EB					92	81	D 2345

The Fast Error report includes the information needed by a Digital Services engineer to troubleshoot a problem with a tape or disk device. ♦

18.5.6 DECevent Restrictions

When you use the DECevent utility, note some of the restrictions listed in this section.

Page File Quota

Sometimes, if the page file quota is exceeded, DECevent will terminate and return you to the system prompt. If this happens, invoke the last command.

Getting Information About the System

18.5 Using the DECevent Event Management Utility (DECevent) (Alpha Only)

Logical File Names

DECevent does not translate as input any logical defined as a search list of file names. For example:

```
$ DEFINE EVENT_LOG DISK1:[EVENTS]EVENT_LOG1.SYS,DISK1:EVENT_LOG.SYS
$ DIAGNOSE/ANALYZE EVENT_LOG
```

```
DECevent T1.0 FT2
```

```
_DIAGNOSE-FAT: Analyze - No files found ' event_log '
```

```
_DIAGNOSE-FAT: An error occurred while executing a command ruleset
```

```
_DIAGNOSE-INF: No Error Messages to send in thread 1
```

Log File Purging

DECevent does not automatically purge log files. Set the version limit on the files and directory to your preference. For example:

```
$ SET FILE/VERSION=3 DIAG_ACTIVITY.LOG
```

System-Initiated Call Logging

When a system running DECevent is shut down and rebooted, DECEVENT\$STARTUP.COM does not define FMGPROFILE logicals. This can interfere with proper logging of system initiated call logging (SICL) due to missing customer profile information in the SICL message text.

Unrecognized Messages

The DIAGNOSE command does not recognize error log messages logged using the \$SNDERR system service.

18.6 Setting Up, Maintaining, and Printing the Operator Log File

The following sections describe the contents of the operator log file and OPCOM messages. They also explain how to perform the following tasks, which require OPER privilege:

Task	Section
Setting up the operator log file	Section 18.6.3
Maintaining the operator log file	Section 18.6.4
Printing the operator log file	Section 18.6.5

18.6.1 Understanding the Operator Log File

The operator log file (SYS\$MANAGER:OPERATOR.LOG) records system events and user requests that the Operator Communication Manager (OPCOM) sends to the operator terminal. This recording occurs even if all operator terminals have been disabled. By default, OPCOM starts when you boot your system. (For more information on OPCOM, see Section 2.4.)

You can use the operator log file to anticipate and prevent hardware and software failures and to monitor user requests for disk and magnetic tape operations. By regularly examining the operator log file, you can often detect potential problems and take corrective action.

The size of and access to the OPERATOR.LOG file (or to the file pointed to by the logical OPC\$LOGFILE_NAME) is limited by the size and access of the disk device on which it resides. If disk device does not have enough room to write to the log file or if access to the device in any other way is restricted, records might be missing from the log file.

Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

18.6.2 Understanding OPCOM Messages

The following sections describe the types of messages that appear in the operator log file.

Type of Message	Section
Initialization messages	Section 18.6.2.1
Device status messages	Section 18.6.2.2
Terminal enable and disable messages	Section 18.6.2.3
User request and operator reply messages	Section 18.6.2.4
Volume mount and dismount messages	Section 18.6.2.5
System parameter messages	Section 18.6.2.6
Security alarm messages	Section 18.6.2.7

Section 18.6.2.8 contains an example of typical kinds of messages found in an operator log file.

18.6.2.1 Initialization Messages

When you enter the **REPLY/LOG** command, the system closes the current operator log file and creates and opens a new version of the file. The system records all subsequent OPCOM messages in the new log file.

When you create a new log file, the first message recorded in it is an initialization message. This message shows the terminal name of the operator who initialized the log file, and the log file specification. This message appears in the following format:

```
%%%%%%%%%%  %OPCOM, <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%%
Logfile has been initialized by operator <terminal-name>
Logfile is <logfile-specification>
```

For example:

```
%%%%%%%%%%  OPCOM, 19-APR-1995 12:29:24.52 %%%%%%%%%%%
Logfile has been initialized by operator _MARS$VTA2:
Logfile is HOMER::SYS$SYSMOND:[SYSMGT]OPERATOR.LOG;43
```

18.6.2.2 Device Status Messages

Some I/O drivers send messages to OPCOM concerning changes in the status of the devices they control. For example, when a line printer goes off line, an OPCOM message appears in the operator log file at periodic intervals until you explicitly return the device to online status.

The device status message appears in the operator log file in the following format:

```
%%%%%%%%%%  OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%%
Device <device-name> is offline
```

This message can appear for card readers, line printers, and magnetic tapes.

18.6.2.3 Terminal Enable and Disable Messages

Following are explanations of commands you can give to enable and disable terminals as operator terminals (or consoles) and explanations of the corresponding messages that appear in the operator log file.

Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

REPLY/ENABLE Messages

To designate a terminal as an operator terminal, enter the REPLY/ENABLE command from the desired terminal. OPCOM confirms the request by displaying messages in the following format at the operator terminal and in the operator log file:

```
%%%%%%%%%% %OPCOM dd-mmm-yyyy hh:mm:ss.cc %%%%%%%%%%
Operator <terminal-name> has been enabled, username <user-name>

%%%%%%%%%% %OPCOM dd-mmm-yyyy hh:mm:ss.cc %%%%%%%%%%
Operator status for operator <terminal-name>
<status-report>
```

These messages tell you which terminal has been established as an operator terminal and lists the requests the terminal can receive and respond to.

You can also designate a terminal as an operator terminal for a particular function by entering the REPLY/ENABLE=*class* command.

If you enter the command REPLY/ENABLE=TAPES, for example, OPCOM displays messages similar to the following:

```
%%%%%%%%%% %OPCOM 19-APR-1995 10:25:35.74 %%%%%%%%%%
Operator _ROUND$OPA1: has been enabled, username SYSTEM

%%%%%%%%%% %OPCOM 19-APR-1995 10:25:38.82 %%%%%%%%%%
Operator status for operator _ROUND$OPA1:
TAPES
```

OPCOM confirms that the terminal is established as an operator terminal and indicates that the terminal can only receive and respond to requests concerning magnetic-tape-oriented events, such as the mounting and dismounting of tapes.

REPLY/DISABLE Messages

A terminal that you designate as an operator terminal automatically returns to nonoperator status when the operator logs out. To return the terminal to normal (nonoperator) status without logging out, enter the REPLY/DISABLE command from the terminal.

OPCOM confirms that the terminal is no longer an operator terminal by displaying a message both at the operator terminal and in the operator log file. The message, which tells you which terminal has been restored to nonoperator status and when the transition occurred, has the following format:

```
%%%%%%%%%% %OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%
Operator <terminal-name> has been disabled, username <user-name>
```

If you designate a terminal as an operator terminal and only partial operator status is disabled, OPCOM displays a status message. This message lists which requests the terminal can still receive and respond to. This message is displayed at the operator terminal and in the operator log file in the following format:

```
%%%%%%%%%% %OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%
Operator status for operator <terminal-name>
<status-report>
```

For example, suppose you designate a terminal as an operator terminal that receives messages concerning magnetic tapes and disks, as well as messages intended for the special site-specific operator class known as OPER10. Later, you relinquish the terminal's ability to receive messages concerning tapes. When you enter the REPLY/DISABLE=TAPES command, OPCOM returns a message like the following:

Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

```
%%%%%%%%%% %Opcom 19-APR-1995 09:23:45.32 %%%%%%%%%%
Operator status for operator TTA3
DISKS, OPER10
```

This message tells you that terminal TTA3 still receives and can respond to messages about disks and messages directed to OPER10.

18.6.2.4 User Request and Operator Reply Messages

To communicate with the operator, the user enters the REQUEST command, specifying either the /REPLY or /TO qualifier. Following are explanations of these qualifiers:

Command	Explanation
REQUEST /REPLY	<p>If the user enters this command, the request is recorded in the operator log file in the following format:</p> <pre>%%%%%%%%%% %OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%% Request <request-id>, from user <user-name> on <node-name> <_terminal-name>:, <"message-text"></pre> <p>This message tells you which user sent the message, the time the message was sent, the request identification number assigned to the message, the originating terminal, and the message itself.</p>
REQUEST /TO	<p>If the user enters this command, the request is recorded in the operator log file in the format shown in the REQUEST/REPLY example, but without a request identification number:</p> <pre>%%%%%%%%%% %OPCOM, <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%% Request from user <user-name> on <node-name> <_terminal-name>:, <"message-text"></pre>

Messages also differ depending on how you reply to a user:

Command	Explanation
REPLY/TO	<p>When you respond to a user's request and specify the /TO qualifier, the response is recorded in the operator log file in the following format:</p> <pre>response message <hh:mm:ss.cc>, request <request-id> completed by operator <terminal-name></pre> <p>This message indicates how the operator responded to the user's request, as well as when the response was entered and which operator responded.</p>
REPLY /ABORT	<p>When you respond to a user's request and specify the /ABORT qualifier, the response is recorded in the operator log file in the following format:</p> <pre><hh:mm:ss.cc>, request <request-id> was aborted by operator <terminal-name></pre>
REPLY /PENDING	<p>When you respond to a user's request using the /PENDING qualifier, the response is not recorded in the operator log file because the request has not yet been completed (that is, the request has not been fulfilled or aborted).</p>

When a user enters a REQUEST/REPLY command and you have disabled all terminals as operators' terminals, OPCOM records all subsequent users' requests in the log file, but returns a message to the user indicating that no operator coverage is available.

All other OPCOM responses to REPLY commands, except responses involving the REPLY/ENABLE, REPLY/DISABLE, and REPLY/LOG commands, are not logged in the operator log file.

Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

18.6.2.5 Volume Mount and Dismount Messages

Perhaps the widest range of operator messages occurs with volume mounts and dismounts; for example:

```
%%%%%%%%%% OPCOM, 19-APR-1995 22:41:07.54 %%%%%%%%%%
message from user SYSTEM
Volume "KLATU" dismounted, on physical device MTA0:
15-APR-1995 22:42:14.81, request 2 completed by operator OPA0
```

18.6.2.6 System Parameter Messages

Users with the appropriate privileges can change the following sets of values for system parameters:

Values	Description
Current	Values stored in the default parameter file on disk and used to boot the system
Active	Values stored in memory and used while the system is running

When the system boots, it reads the current values into memory, creating active values. An active value remains equal to the current value until you change either value.

Users can make the following changes to active and current system parameters:

- Active system parameters—Users with CMKRNL privilege can use the System Management utility (SYSMAN) or the System Generation utility (SYSGEN) to change system parameters in the running (active) system. Users can change only those active values that are categorized as *dynamic* system parameters.
- Current system parameters—Users with SYSPRV privilege can use SYSMAN or SYSGEN to change system parameters in the current system.

Note

Digital recommends that you use AUTOGEN or SYSMAN, not SYSGEN, to change system parameters, as explained in Section 14.2.

OPCOM logs all changes made to active and current system parameters with messages in the following format:

```
%%%%%%%%%% %OPCOM <dd-mm-yy hh:mm:ss.cc> %%%%%%%%%%
Message from user <user-name>
%SYSGEN-I-WRITExxx, <system-mode> system parameters modified by process ID
<process-id> into file <file-spec>
```

For example:

```
%%%%%%%%%% %OPCOM 3-JUN-1995 08:11:59.55 %%%%%%%%%%
Message from user D_PLUTO on ANASAT
%SYSGEN-I-WRITECUR, CURRENT system parameters modified by process ID 0000020B
into file SYS$UPDATE:[SYSTEM]UPDATESYS.PAR;2
```


Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

This message indicates that current system parameters have been changed.

Note

If you have changed the format of system messages with the DCL command SET MESSAGE, these messages might not appear in the log file.

18.6.2.7 Security Alarm Messages

Alarm messages are sent to the security operator terminal when selected events occur. See Section 18.7.6 for instructions about how to enable a terminal to receive security alarm messages.

The following example shows a security alarm OPCOM message after a change to JTQUOTA:

```
%%%%%%%% OPCOM 6-JAN-1995 10:41:21.10 %%%%%%%%%
Message from user AUDIT$SERVER on BISCO
Security alarm (SECURITY) and security audit (SECURITY) on BISCO, system id:
20353
Auditable event:      System UAF record modification
Event time:           6-JAN-1995 10:41:20.69
PID:                  00600123
Process name:         SYSTEM
Username:             SYSTEM
Process owner:        [SYSTEM]
Terminal name:        RTA1:
Image name:           BISCO$DUA0:[SYS0.SYSCOMMON.][SYSEXE]AUTHORIZE.EXE
Object class name:    FILE
Object name:          SYSS$SYSTEM:SYSUAF.DAT;4
User record:          NEWPORT
JTQUOTA:              New:          2048
                     Original:      1024
```

18.6.2.8 Contents of an Operator Log File

Example 18-6 illustrates some typical messages found in an operator log file.

Example 18-6 Sample Operator Log File (SYSS\$MANAGER:OPERATOR.LOG)

```
%%%%%%%% OPCOM, 19-APR-1995 22:26:07.90 %%%%%%%%%
1 Device DMA0: is offline.
Mount verification in progress.
%%%%%%%% OPCOM, 19-APR-1995 22:26:20.22 %%%%%%%%%
Mount verification completed for device DMA0:
%%%%%%%% OPCOM, 19-APR-1995 22:33:54.07 %%%%%%%%%
2 Operator '_ZEUS$VT333:' has been disabled, user JONES
%%%%%%%% OPCOM, 19-APR-1995 22:34:15.47 %%%%%%%%%
Operator '_ZEUS$VT333:' has been enabled, user SMITH
%%%%%%%% OPCOM, 19-APR-1995 22:34:15.57 %%%%%%%%%
operator status for '_ZEUS$VT333:'
PRINTER, TAPES, DISKS, DEVICES
%%%%%%%% OPCOM, 19-APR-1995 22:38:53.21 %%%%%%%%%
3 request 1, from user PUBLIC
Please mount volume KLATU in device MTA0:
The tape is in cabinet A
```

(continued on next page)

Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

Example 18-6 (Cont.) Sample Operator Log File (SYS\$MANAGER:OPERATOR.LOG)

```
##### OPCOM, 19-APR-1995 22:39:54.37 #####
request 1 was satisfied.
##### OPCOM, 19-APR-1995 22:40:23.54 #####
4 message from user SYSTEM
Volume "KLATU" " mounted, on physical device MTA0:
##### OPCOM, 19-APR-1995 22:40:38.02 #####
request 2, from user PUBLIC
MOUNT new relative volume 2 () on MTA0:
##### OPCOM, 19-APR-1995 22:41:07.54 #####
message from user SYSTEM
Volume "KLATU" " dismounted, on physical device MTA0:
15-APR-1995 22:42:14.81, request 2 completed by operator OPA0
##### OPCOM, 19-APR-1995 22:46:47.96 #####
request 4, from user PUBLIC
_TTB5:, This is a sample user request with reply expected.
##### OPCOM, 19-APR-1995 22:47:38.50 #####
request 4 was canceled
##### OPCOM, 19-APR-1995 22:48:21.15 #####
message from user PUBLIC
_TTB5:, This is a sample user request without a reply expected.
##### OPCOM, 19-APR-1995 22:49:37.64 #####
Device DMA0: has been write locked.
Mount verification in progress.
##### OPCOM, 19-APR-1995 23:33:54.07 #####
message from user NETACP
DECnet shutting down
```

The following messages appear in the example:

- ❶ Device status message (see Section 18.6.2.2)
- ❷ Terminal enable and disable message (see Section 18.6.2.3)
- ❸ User request and operator reply messages (see Section 18.6.2.4)
- ❹ Volume mount and dismount messages (see Section 18.6.2.5)

18.6.3 Setting Up the Operator Log File

The operator log file normally resides on the system disk in the [SYSMGR] directory. You can, however, maintain the log file in a different location by defining the logical name OPC\$LOGFILE_NAME.

Because this file is in ASCII format, you can print it. Print copies regularly and retain these copies for reference. Section 18.6.5 describes how to print copies of the operator log file.

The system creates a new version of OPERATOR.LOG each time the system is rebooted (except on workstations in a VMScluster environment, where the log file is not opened by default). Note that one operator log file exists for each node; it is not a shared file.

Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

18.6.3.1 Creating a New Version of the Operator Log File

You can use the DCL command `REPLY/LOG` to create a new version of the file at any time. The highest version is always the one in use and is inaccessible to other users. By default, messages of all operator classes are in the log file.

Following are guidelines for using the `REPLY/LOG` command:

- You can use the `REPLY/LOG/ENABLE=(list-of-classes)` and `REPLY/LOG/DISABLE=(list-of-classes)` commands to specify which operator classes to include in the log file.
- When you use the `/LOG` qualifier with the `REPLY/ENABLE` and `REPLY/DISABLE` commands, the classes you select are enabled or disabled for the log file rather than for the terminal.

If a log file is already open, the list of classes is preserved and enabled on the newly created log file. If a log file is not open, the value of the logical `OPC$ENABLE_LOGFILE_CLASSES` is used. If that logical does not exist, all classes are enabled on the new log file.

For more information, see the `REPLY/LOG`, `REPLY/ENABLE`, and `REPLY/DISABLE` commands in the *OpenVMS DCL Dictionary*.

Example

The following command opens a log file to include messages about mounting and dismounting disks and tapes:

```
$ REPLY/LOG/ENABLE=(DISKS,TAPES)
```

18.6.3.2 Specifying Logical Names

You can specify the default state of the operator log files by defining logical names in the command procedure `SYS$MANAGER:SYLOGICALS.COM`. The following table lists these logical names and their functions. For more information on `SYLOGICALS.COM`, see Section 5.2.5.

Logical Name	Function
<code>OPC\$LOGFILE_ENABLE</code>	Specifies whether an operator log file is opened. If defined to be true, an operator log file is opened. If defined to be false, no operator log file is opened. By default, a log file is opened on all systems except workstations in a VMScluster environment.
<code>OPC\$LOGFILE_CLASSES</code>	Specifies the operator classes that are enabled for the log file. By default, a log file is opened for all classes. The logical name can be a search list of the allowed classes, a comma-separated list, or a combination of the two. Note that you can define <code>OPC\$LOGFILE_CLASSES</code> even if you do not define <code>OPC\$LOGFILE_ENABLE</code> . In that case, the classes are used for any log files that are opened, but the default is used to determine whether to open the log file.
<code>OPC\$LOGFILE_NAME</code>	Specifies the name of the log file. By default, the log file is named <code>SYS\$MANAGER:OPERATOR.LOG</code> . If you specify a disk other than the system disk, include commands to mount that disk in the command procedure <code>SYLOGICALS.COM</code> .
<code>OPC\$OPA0_ENABLE</code>	Overrides values of symbols for workstations in a cluster. If you define the logical as <code>TRUE</code> , it sets the <code>OPA0</code> device to <code>BROADCAST</code> (overrides the <code>NOBROADCAST</code> default setting). For systems that are not workstations in a cluster, if you define the logical as <code>FALSE</code> , it sets the <code>OPA0</code> device to <code>NOBROADCAST</code> .

Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

Note

The only logical that is used for more than the initial startup of OPCOM is OPC\$LOGFILE_NAME. All other OPCOM logicals are ignored. For example, a REPLY/LOG command opens a new operator log file even if the logical OPC\$LOGFILE_ENABLE is defined to be false. To reset OPCOM states and classes after startup, use REPLY/ENABLE or REPLY/DISABLE commands.

18.6.4 Maintaining the Operator Log File

Devise a plan for regular maintenance of operator log files. One way is to start a new log file and rename the second-highest version daily. (See the example in the next section.) You might want to purge outdated versions of the operator log file on a regular basis. However, do not delete versions that you have not backed up. For more information, see Section 5.2.7.9.

If OPCOM is inadvertently deleted, follow these steps to start it manually:

1. Log in to the SYSTEM account so that you have the required privileges to perform the operation.
2. Enter the following command to execute the startup command procedure (STARTUP.COM), specifying OPCOM as the command parameter:

```
$ @SYS$SYSTEM:STARTUP OPCOM
```

18.6.5 Printing the Operator Log File

Perform the following operation to produce a printed copy of the most recent version of the operator log file. (You must have OPER privilege.)

1. Use the following command to enable the terminal as an operator terminal:

```
$ REPLY/ENABLE
```

2. Close the current log file and open a new one by entering the following command:

```
$ REPLY/LOG
```

3. Set the default to SYS\$MANAGER and enter the following command to list all versions of the file:

```
$ SET DEFAULT SYS$MANAGER  
$ DIRECTORY OPERATOR.LOG
```

4. Rename the second-highest version to OPERATOR.OLD:

```
$ RENAME OPERATOR.LOG;-1 OPERATOR.OLD
```

The version number, -1, specifies that you want to rename the second-highest version of this file. (The highest version number is the current operator log file.)

5. Print the operator log file by entering the following command:

```
$ PRINT OPERATOR.OLD
```


Getting Information About the System

18.6 Setting Up, Maintaining, and Printing the Operator Log File

Example

```
❶ $ REPLY/ENABLE
❷ $ REPLY/LOG

##### OPCOM, 19-APR-1995 12:28:20.11 #####
❸ Logfile was closed by operator _MARS$VTA2:
Logfile was HOMER::SYS$MANAGER:[SYSMGT]OPERATOR.LOG;27
##### OPCOM, 19-APR-1995 12:29:24.52 #####
Logfile has been initialized by operator _MARS$VTA2:
Logfile is HOMER::SYS$MANAGER:[SYSMGT]OPERATOR.LOG;28

❹ $ SET DEFAULT SYS$MANAGER
❺ $ DIRECTORY OPERATOR.LOG

Directory SYS$MANAGER:[SYSMGT]
OPERATOR.LOG;28          OPERATOR.LOG;27

Total of 2 files.

❻ $ RENAME OPERATOR.LOG;-1 OPERATOR.OLD
❼ $ PRINT OPERATOR.OLD
```

Following are explanations of the numbered commands and responses in this example:

- ❶ The REPLY/ENABLE command enables the terminal as an operator terminal.
- ❷ The REPLY/LOG command closes the current log file and opens a new one.
- ❸ The response from OPCOM verifies that it has opened a new log file.
- ❹ The SET DEFAULT command sets the operator default disk to the system disk.
- ❺ The DIRECTORY command displays the files in the directory [SYSMGT] on the system disk.
- ❻ The RENAME command renames the second-highest version of the operator log file to OPERATOR.OLD.
- ❼ The PRINT command prints the old operator log file, OPERATOR.OLD.

18.7 Using Security Auditing

This section discusses how security auditing works; it also explains how to enable security auditing and how to create a new version of the security audit log file. For more information about the security audit log file, see the *OpenVMS Guide to System Security*.

18.7.1 Understanding Security Auditing

Security auditing is the act of recording security-relevant events as they occur on the system. Security-relevant events are divided into a number of categories called **event classes**.

By default, the system enables security auditing when you install or upgrade your system for the events shown in Table 18–6.

Getting Information About the System

18.7 Using Security Auditing

Table 18–6 Event Classes Audited by Default

Class	Description
ACL	Access to any object holding a security Auditing ACE.
Audit	All uses of the SET AUDIT command. You cannot disable this category.
Authorization	All changes to the authorization database: <ul style="list-style-type: none">• System user authorization file (SYSUAF)• Network proxy authorization files: NETPROXY and †NET\$PROXY• Rights database (RIGHTSLIST)
Breakin	All break-in attempts: batch, detached, dialup, local, network, remote.
Logfailure	All login failures: batch, dialup, local, remote, network, subprocess, detached.
†VAX specific	

If the security requirements at your site justify additional auditing, you can enable security auditing for other event classes by using the DCL command SET AUDIT, as explained in Section 18.7.4.

18.7.1.1 Security Audit Log File

The audit server process, which is created at system startup, records the events that are shown in Table 18–6 in the security audit log file, SYS\$MANAGER:SECURITY.AUDIT\$JOURNAL.

The usefulness of the security audit log file depends upon the procedures you adopt to review the file on a regular basis. For example, you might implement the following procedure as part of your site audit review policy:

1. Create a new version of the security audit log file each morning.
2. Review the previous version of the log file for suspicious system activity. Depending on the number of security events you are auditing on your system, it might be impractical to review every audit record written to the audit log file. In that case, you might want to select a specific set of records from the log file (for example, all Authorization and Breakin records, or all events created outside normal business hours).
3. If, during your review, you find any security events that appear suspicious, perform a more detailed inspection of the security audit log file, as described in the *OpenVMS Guide to System Security*.

18.7.1.2 Audit Log Files in Mixed-Version Clusters

The Audit Analysis utility (ANALYZE/AUDIT) running on Version 1.*n* systems is unable to process the current version of audit log files. You must use the current version of ANALYZE/AUDIT to process the current version of the audit log files. The recommended procedure is to maintain separate audit log files on mixed-version clusters.

Getting Information About the System

18.7 Using Security Auditing

If redirecting the audit log files, issue the following command on both a Version 1.n node and on the node running the current version:

```
AUDIT/JOURNAL/DESTINATION=filespec
```

The destination filespec is stored in the audit server database file. By default, the files are stored in SYS\$COMMON:[SYSMGR] and are called SECURITY_AUDIT.AUDIT\$JOURNAL and SECURITY.AUDIT\$JOURNAL, respectively. See the *OpenVMS Guide to System Security* for further information.

18.7.2 Displaying Security Auditing Information

To see which event classes your site currently audits, you can enter the DCL command SHOW AUDIT.

Following is an example of security information:

```
$ SHOW AUDIT
```

System security alarms currently enabled for:

ACL

Breakin: dialup,local,remote,network,detached

Privilege use:

SECURITY

Privilege failure:

SECURITY

System security audits currently enabled for:

ACL

Authorization

Breakin: dialup,local,remote,network,detached

Login: dialup,local,remote,network,detached

Logfailure: batch,dialup,local,remote,network,subprocess,detached

Logout: dialup,local,remote,network,detached

Privilege use:

SECURITY

Privilege failure:

ACNT	ALLSPOOL	ALTPRI	AUDIT	BUGCHK	BYPASS	CMEXEC	CMKRNL
DETACH	DIAGNOSE	EXQUOTA	GROUP	GRPNAM	GRPPRV	LOG_IO	MOUNT
NETMBX	OPER	PFNMAP	PHY_IO	PRMCEB	PRMGBL	PRMMBX	PSWAPM
READALL	SECURITY	SETPRV	SHARE	SHMEM	SYSGBL	SYSLCK	SYSNAM
SYSPRV	TMPMBX	VOLPRO	WORLD				

DEVICE access:

Failure: read,write,physical,logical,control

FILE access:

Failure: read,write,execute,delete,control

VOLUME access:

Failure: read,write,create,delete,control

18.7.3 Delaying Startup of Auditing

Ordinarily, the system turns on auditing in VMS\$LPBEGIN just before SYSTARTUP_VMS.COM executes. You can change this behavior, however, by redefining the logical name SYS\$AUDIT_SERVER_INHIBIT.

To change the point at which the operating system begins to deliver security-event messages, add the following line to the SYS\$MANAGER:SYLOGICALS.COM command procedure:

```
$ DEFINE/SYSTEM/EXECUTIVE SYS$AUDIT_SERVER_INHIBIT YES
```


Getting Information About the System

18.7 Using Security Auditing

You can initiate auditing during another phase of system startup, perhaps at the end of SYSTARTUP_VMS.COM, by editing the command file to add the following line:

```
$ SET AUDIT/SERVER=INITIATE
```

For information on editing SYSTARTUP_VMS.COM, see Section 5.2.7.

18.7.4 Enabling Security Auditing for Additional Classes

To enable security auditing for classes in addition to those shown in Table 18–6, use the following format:

```
SET AUDIT/ENABLE=event-class[...] {/ALARM | /AUDIT}
```

The *OpenVMS Guide to System Security* contains descriptions of event classes that you can enable.

When you enable auditing for additional event classes, you must specify two qualifiers:

1. /ENABLE
2. either /ALARM or /AUDIT (Although you *must* specify one qualifier, you *can* specify both.)

Following are explanations of the /ENABLE, /ALARM, and /AUDIT qualifiers.

Qualifier	Explanation
/ENABLE	Defines which event classes you want audited. See Chapter 19 for more information.
/ALARM /AUDIT	<p>Defines the destination of the event message.</p> <ul style="list-style-type: none">• /ALARM directs the message to all enabled security operator terminals.• /AUDIT directs the message to the security audit log file. <p>Use the /ALARM and /AUDIT qualifiers to report critical events. Less critical events can be written only to the security audit log file for later examination.</p> <p>The default event classes listed in Table 18–6 are sent as both alarms and audits.</p>

The system begins auditing new events on all nodes as soon as you enable them.

Examples

1. The command in the following example enables auditing for volume mounts and dismounts and sends messages to the security audit log file.

```
$ SET AUDIT/ENABLE=MOUNT/AUDIT
```

2. The command in the following example enables auditing of unsuccessful file accesses and sends messages to all enabled security operator terminals as well as to the security audit log file.

```
$ SET AUDIT/ALARM/AUDIT/ENABLE=ACCESS=FAILURE/CLASS=FILE
```


18.7.5 Disabling Security Auditing

The system continues auditing until you explicitly disable the classes with the /DISABLE qualifier using the following syntax:

```
SET AUDIT/DISABLE=event-class[...] [/ALARM | /AUDIT]
```

18.7.6 Enabling a Terminal to Receive Alarm Messages

The system sends alarm messages to terminals enabled for security class messages. Security alarm messages are not written to the operator log file. They appear only on terminals enabled for security class messages.

In most cases, security alarm messages appear on the system console by default. Since messages scroll quickly off the screen, it is good practice to enable a separate terminal for security class messages and disable message delivery to the system console.

Either choose a terminal in a secure location that provides hardcopy output, or have dedicated staff to monitor the security operator terminal. You can enable any number of terminals as security operators.

To set up a terminal to receive security class alarms, enter the following DCL command from the designated terminal:

```
$ REPLY/ENABLE=SECURITY
```

The following example shows a security alarm message:

```
***** OPCOM 25-MAY-1995 16:07:09.20 *****
Message from user AUDIT$SERVER on GILMORE
Security alarm (SECURITY) on GILMORE, system id: 20300
Auditable event:      Process suspended ($SUSPND)
Event time:           25-MAY-1995 16:07:08.77
PID:                  30C00119
Process name:         Hobbit
Username:             HUBERT
Process owner:        [LEGAL,HUBERT]
Terminal name:        RTA1:
Image name:           $99$DUA0:[SYS0.SYSCOMMON.][SYSEXEC]SET.EXE
Status:               %SYSTEM-S-NORMAL, normal successful completion
Target PID:           30C00126
Target process name:   SMISERVER
Target username:      SYSTEM
Target process owner:  [SYSTEM]
```

18.7.7 Generating Security Reports

The most common type of report to generate is a brief, daily listing of events. You can create a command procedure that runs in a batch job every evening before midnight to generate a report of the day's security event messages and send it to the system manager via MAIL.

Note

Since the MOUNT command translates /NOLABEL to /FOREIGN in the audit record, use ANALYZE/AUDIT/SELECT=MOUNT_FLAGS=FOREIGN instead of ANALYZE/AUDIT/SELECT=MOUNT_FLAGS=NOLABEL.

Getting Information About the System

18.7 Using Security Auditing

The following example shows the ANALYZE/AUDIT command line you would use to generate this type of report:

```
$ ANALYZE/AUDIT/SINCE=TODAY/OUTPUT=31JAN1995.AUDIT -  
_ $ SYS$MANAGER:SECURITY.AUDIT$JOURNAL  
$ MAIL/SUBJECT="Security Events" 31JAN1995.AUDIT SYSTEM
```

18.7.8 Creating a New Version of the Security Audit Log File

Because the security audit log file continues to grow until you take action, you must devise a plan for maintaining it.

You use the following SET AUDIT command to create a new version of the clusterwide security audit log file. To prevent the loss of audit messages, the previous version of the audit log file is not closed until all audit messages stored in memory are written to the file.

18.7.8.1 Creating a New Clusterwide Version of the Log File

To open a new, clusterwide version of the security audit log file, use the following command:

```
$ SET AUDIT/SERVER=NEW_LOG
```

The audit server process opens a new version of the audit log file on each cluster node.

After you open the new log, rename the old version, using a naming convention for your files that incorporates in the file name a beginning or ending date for the data. Then copy the file to another disk, delete the log from the system disk to save space, and run the Audit Analysis utility on the old log.

By archiving this file, you maintain a clusterwide history of auditing messages. If you ever discover a security threat on the system, you can analyze the archived log files for a trail of suspicious user activity during a specified period of time.

18.7.8.2 Creating a New Node-Specific Version of the Log File

In some cases, VMScluster nodes might not share the same system security audit log file. To create a new, node-specific version of the security audit log file, use the following commands:

```
$ SET AUDIT/DESTINATION=filespec  
$ SET AUDIT/SERVER=NEW_LOG
```

For the *filespec*, include a logical name that points to a node-specific file; for example, SYS\$SPECIFIC:[SYSMGR]SECURITY. System security audit log files on other nodes are unaffected.

18.8 Monitoring Operating System Performance

The Monitor utility (MONITOR) is a system management tool that you can use to obtain information on operating system performance. Various MONITOR qualifiers collect system performance data from the running system or play back data recorded previously in a recording file. When you play back data, you can display it, summarize it, and even rerecord it to reduce the amount of data in the recording file.

Following an explanation of the Monitor utility are sections that tell how to perform these tasks:

Getting Information About the System

18.8 Monitoring Operating System Performance

Task	Section
Invoking the Monitor utility	Section 18.8.2
Using live display monitoring	Section 18.8.3
Using live recording monitoring	Section 18.8.4
Using concurrent display and recording monitoring	Section 18.8.5
Using playback monitoring	Section 18.8.6
Using remote playback monitoring	Section 18.8.7
Rerecording monitoring	Section 18.8.8
Running MONITOR continuously	Section 18.8.9

For additional information about interpreting the information the Monitor utility provides, see the *Guide to OpenVMS Performance Management*. For additional information about using the Monitor utility, see the *OpenVMS System Management Utilities Reference Manual*.

18.8.1 Understanding the Monitor Utility (MONITOR)

Using MONITOR, you can monitor classes of systemwide performance data (such as system I/O statistics, page management statistics, and time spent in each of the processor modes) at specifiable intervals, and produce several types of output. You can also develop a database of performance information for your system by running MONITOR continuously as a background process, as explained in Section 18.8.9.

18.8.1.1 MONITOR Classes

Each MONITOR class consists of data items that, taken together, provide a statistical measure of a particular system performance category. The data items defined for individual classes are listed in the description of the MONITOR command in the *OpenVMS System Management Utilities Reference Manual*.

To monitor a particular class of information, you specify a class name on the MONITOR command line. The information MONITOR displays depends on the type of class you select. Table 18–7 compares the two MONITOR class types.

Table 18–7 Types of MONITOR Classes

Type of class	Description
System	Provides statistics on resource use for the entire system
Component	Provides statistics on the contribution of individual components to the overall system or VMScluster

As an example of the distinction between types of MONITOR classes, the IO class includes a data item to measure all direct I/O operations for the entire system, and is therefore a system class. The DISK class measures direct I/O operations for individual disks, and is therefore a component class.

Table 18–8 describes each MONITOR class and indicates whether it is a system or component class.

Getting Information About the System

18.8 Monitoring Operating System Performance

Table 18–8 MONITOR Classes

Class	Type	Description
ALL_CLASSES	System or Component	Statistics for all classes
CLUSTER	System	Clusterwide performance statistics
DECNET	System	DECnet for OpenVMS statistics
DISK	Component	Disk I/O statistics
DLOCK	System	Distributed lock management statistics
FCP	System	File control primitive statistics
FILE_SYSTEM_CACHE	System	File system cache statistics
IO	System	System I/O statistics
LOCK	System	Lock management statistics
MODES	Component	Time spent in each of the processor modes
MSCP_SERVER	System	MSCP server statistics
PAGE	System	Page management statistics
PROCESSES	Component	Statistics on all processes
RMS	Component	Record Management Services statistics
SCS	Component	System Communications Services statistics
STATES	System	Number of processes in each of the scheduler states
SYSTEM	System	Summary of statistics from other classes
TRANSACTION	System	DECdtm services statistics
†VBS	System	Virtual balance slot statistics
VECTOR	System	Vector processor scheduled usage
†VAX specific		

18.8.1.2 Display Data

Except in the PROCESSES class, all data item statistics are displayed as rates or levels:

- **Rates** are shown in number of occurrences per second.
- **Levels** are values that indicate the size of the monitored data item.

You can request any or all of four different statistics for each data item:

Statistic	Description
Current rate or level	Most recently collected value for the rate or level
Average rate or level	Measured from the beginning of the MONITOR request
Minimum rate or level	Measured from the beginning of the MONITOR request
Maximum rate or level	Measured from the beginning of the MONITOR request

For the DISK, MODES, SCS, and STATES classes, you can optionally express all statistics as percentages.

In the PROCESSES class, MONITOR displays descriptive information, level information, and counters that increase over time.

Getting Information About the System

18.8 Monitoring Operating System Performance

18.8.1.3 Output Types

MONITOR collects system performance data by class and produces three forms of optional output, depending on the qualifier you specify:

Qualifier	Description
/DISPLAY	Produces output in the form of ASCII screen images, which are written at a frequency governed by the /VIEWING_TIME qualifier.
/RECORD	Produces a binary recording file containing data collected for requested classes; one record for each class is written per interval.
/SUMMARY	Produces an ASCII file containing summary statistics for all requested classes over the duration of the MONITOR request.

If you specify /INPUT with any of these qualifiers, MONITOR collects performance data from one or more previously created recording files; otherwise, data is collected from counters and data structures on the running system.

You use the /BEGINNING and /ENDING qualifiers to specify, respectively, when you want a MONITOR request to begin and end.

Using the /DISPLAY Qualifier

Information collected by MONITOR is normally displayed as ASCII screen images. You can use the optional /DISPLAY qualifier to specify a disk file to contain the information. If you omit the file specification, output is directed to SYS\$OUTPUT.

Note

Be careful when you use the /DISPLAY qualifier. Because MONITOR enters display information into the file continuously, its size can grow very quickly.

See the *OpenVMS System Management Utilities Reference Manual* for a discussion of the /DISPLAY qualifier.

Using the /RECORD Qualifier

When you use the /RECORD qualifier, all data pertaining to the class is recorded, even if you are concurrently displaying only a single statistic or a single item of a component statistics class. The file is created when a MONITOR request is initiated and closed when a request terminates. You can use the resulting file as a source file for later requests to format and display the data on a terminal, to create a summary file, or to create a new recording file with different characteristics.

18.8.2 Invoking the Monitor Utility

To invoke the Monitor utility, enter the following DCL command:

```
$ MONITOR
```

MONITOR then displays the following prompt:

```
MONITOR>
```

In response to the prompt, you can enter any of the MONITOR commands, which are described in *OpenVMS System Management Utilities Reference Manual*. The most frequently used MONITOR command, however, specifies a class name.

Getting Information About the System

18.8 Monitoring Operating System Performance

Example

MONITOR> MONITOR PAGE

In this example, you specify the PAGE class name in the MONITOR command to monitor page management statistics.

You can also use the MONITOR command from DCL command level.

How to Override or Terminate a MONITOR Request

Generally, each MONITOR request runs until the time specified or implied by the /ENDING qualifier. However, to override or terminate a MONITOR request, you can press one of the following:

Keys	Description
Ctrl/W	Temporarily overrides a /VIEWING_TIME value and generates a new display immediately following the current one. This feature is useful when a broadcast message overwrites the MONITOR display area. You can also use Ctrl/W in conjunction with a large /VIEWING_TIME value to generate display events on demand.
Ctrl/C	Terminates the current request without exiting from the utility. You can then initiate a new request or enter any MONITOR command at the MONITOR> prompt.
Ctrl/Z	Terminates the current request and exits from MONITOR.

18.8.3 Using Live Display Monitoring

Use the live display monitoring mode of operation when you want to examine the activity of a running system, either on a routine basis or as part of an installation checkout, tuning, or troubleshooting exercise. The system does not keep a historical record of output. The following examples show how to use the live display monitoring mode.

Examples

1. \$ MONITOR PROCESSES/TOPCPU

The command displays a bar graph showing the eight processes that were the top consumers of CPU time during the period between displays. It also displays the amount of CPU time each process used.

The command might produce a display similar to the following:

```
OpenVMS Monitor Utility
TOP CPU TIME PROCESSES
on node BOMBAY
20-JAN-1995 10:06:49

07E00181 CAFARET 100
0          25      50      75      100
+ - - - - + - - - - + - - - - + - - - - +
*****
|          |          |          |          |
|          |          |          |          |
|          |          |          |          |
|          |          |          |          |
|          |          |          |          |
|          |          |          |          |
+ - - - - + - - - - + - - - - + - - - - +
```


Getting Information About the System

18.8 Monitoring Operating System Performance

This example shows that user CAFARET is using 100 percent of the CPU time available. To display more information about the computer resources a user is using, use a command similar to the following:

```
$ SHOW PROCESS/CONTINUOUS/ID=07E00181
```

For this example, the most useful information in the resulting display is the name of the image at the end of the display; for example:

```
.  
.  
.  
$1SDUA1:[SYS1D.SYSCOMMON.] [SYSEXE] RODAN.EXE
```

This example indicates that CAFARET is running RODAN.EXE, which might be new software that is unintentionally running in a loop. This situation would occur if CAFARET were a privileged user running a process at a higher priority than other users.

2. `$ MONITOR/DISPLAY=PROCESSES.LOG PROCESSES`

You can route MONITOR display output to any supported terminal device or to a disk file. This command writes MONITOR's display process statistics to the file PROCESSES.LOG. You can then print this file on a hardcopy device.

Caution

Because data is continuously added to the display file, be careful that the file does not grow too large.

3.

```
$ MY_CLASSES ::= -  
_ $ "DECNET+FCP+IO+LOCK+MODES+PAGE+PROCESSES+STATES"  
$ MONITOR/NODE=(CURLEY,LARRY)/INTERVAL=20/VIEWING_TIME=8 'MY_CLASSES'
```

You might find it convenient to establish DCL symbols for frequently used combinations of class names, as in this example. The MONITOR command collects selected classes of data for VMScluster nodes CURLEY and LARRY every 20 seconds. Every 8 seconds, the system displays the most recently collected data for one of the classes. MONITOR predetermines the ordering of the classes for display.

18.8.4 Using Live Recording Monitoring

Use live recording to capture MONITOR data for future use. Possible uses include the following:

- Installation checkout, tuning, troubleshooting; that is, all the uses that are listed for live display monitoring.

Choose recording over display when you want to capture more classes than you can reasonably watch at a terminal, when a terminal is not available, or when you want to gather data about the system but cannot spend time at the terminal until later.

- Routine performance data gathering for long-term analysis.

Getting Information About the System

18.8 Monitoring Operating System Performance

You can record MONITOR data on a routine basis and summarize it to gather data about system resource use over long periods of time.

Caution

Because data is continuously added to the recording file, be careful that the file does not grow too large.

The following example shows how to use the live recording mode of operation.

Example

```
$ MONITOR/NODE=(LARRY,MOE)/NODISPLAY/RECORD MODES+STATES
```

The command in this example records data on the time spent in each processor mode and on the number of processes in each scheduler state for nodes LARRY and MOE. The command does not display this information.

18.8.5 Using Concurrent Display and Recording Monitoring

Use the concurrent display and recording mode of operation when you want to both retain performance data and watch as it is being collected. Because MONITOR allows shared read access to the recording file, a separate display process can play back the recording file as it is being written by another process.

The following examples show how to use the concurrent display and recording mode of operation. The first example both collects and records data in the same command. The second and third examples show how you can perform concurrent recording and display using two separate processes: the process in the second example performs recording; the process in the third example plays back the file to obtain a summary.

Examples

1.

```
$ MONITOR/RECORD FCP/AVERAGE,FILE_SYSTEM_CACHE/MINIMUM
```

This command collects and records file system data and file system cache data every 3 seconds. It also displays, in bar graph form, average FCP statistics and minimum FILE_SYSTEM_CACHE statistics. The display alternates between the two graphs every 3 seconds. You can obtain current statistics in a subsequent playback request.

2.

```
$ MONITOR/RECORD=SYS$MANAGER:ARCHIVE.DAT -  
_$/INTERVAL=300/NODISPLAY ALL_CLASSES
```

This command archives data for all classes once every 5 minutes. You might find it convenient to execute a similar command in a batch job, taking care to monitor disk space usage.

3.

```
$ MONITOR/INPUT=SYS$MANAGER:ARCHIVE.DAT: -  
_$/NODISPLAY/SUMMARY/BEGINNING="-1" PAGE,IO
```

The command in this example produces a summary of page and I/O activity that occurred during the previous hour, perhaps as part of an investigation of a reported performance problem. Note that because the recording process executes an OpenVMS RMS flush operation every 5 minutes, up to 5 minutes of the most recently collected data is not available to the display process.

You can specify the time between flush operations explicitly with the /FLUSH_INTERVAL qualifier. Note also that the display process must have read access to the recording file.

18.8.6 Using Playback Monitoring

Use playback of a recording file to obtain terminal output and summary reports of all collected data or a subset of it. You can make a subset of data according to class, node, or time segment. For example, if you collect several classes of data for an entire day, you can examine or summarize the data on one or more classes during any time period in that day.

You can also display or summarize data with a different interval than the one at which it was recorded. You control the actual amount of time between displays of screen images with the `/VIEWING_TIME` qualifier. The following examples show how to use the playback mode of operation.

Examples

1. `$ MONITOR/RECORD/INTERVAL=5 IO`

```
.  
.  
.  
$ MONITOR/INPUT IO
```

The commands in this example produce system I/O statistics. The first command gathers and displays data every 5 seconds, beginning when you enter the command and ending when you press Ctrl/Z. In addition, the first command records binary data in the default output file `MONITOR.DAT`.

The second command plays back the I/O statistics display, using the data in `MONITOR.DAT` for input. The default viewing time for the playback is 3 seconds, but each screen display represents 5 seconds of monitored I/O statistics.

2. `$ MONITOR/RECORD/NODISPLAY -`
 `_ $ /BEGINNING=08:00:00 -`
 `_ $ /ENDING=16:00:00 -`
 `_ $ /INTERVAL=120 DISK`

```
.  
.  
.  
$ MONITOR/INPUT/DISPLAY=HOURLY.LOG/INTERVAL=3600 DISK
```

The sequence of commands in this example illustrates data recording with a relatively small interval and data playback with a relatively large interval. This is useful for producing average, minimum, and maximum statistics that cover a wide range of time, but have greater precision than if they had been gathered using the larger interval.

The first command records data on I/O operations for all disks on the system for the indicated 8-hour period, using an interval of 2 minutes. The second command plays the data back with an interval of 1 hour, storing display output in the file `HOURLY.LOG`. You can then type or print this file to show the cumulative average disk use at each hour throughout the 8-hour period.

Note

The current statistic in `HOURLY.LOG` shows the current data in terms of the original collection interval of 120 seconds, not the new collection interval of 3600 seconds.

Getting Information About the System

18.8 Monitoring Operating System Performance

3. \$ MONITOR/INPUT/NODISPLAY/SUMMARY=DAILY.LOG DISK

The command in this example uses the recording file created in the previous example to produce a one-page summary report file showing statistics for the indicated 8-hour period. The summary report has the same format as a screen display. For example:

```

OpenVMS Monitor Utility
DISK I/O STATISTICS
on node TLC
SUMMARY
From: 25-JAN-1995 08:00:00
To: 25-JAN-1995 16:00:00

I/O Operation Rate
CUR AVE MIN MAX
DSA0: SYSTEM_0 0.53 1.50 0.40 3.88
DSA1: SYSTEM_1 0.00 0.39 0.00 8.38
DSA4: WORK_0 0.00 0.11 0.00 1.29
DSA5: WORK_1 0.03 0.87 0.00 5.95
DSA6: WORK_2 0.03 0.25 0.00 2.69
DSA7: WORK_3 0.04 0.97 0.00 20.33
DSA17: TOM_DISK 0.00 0.04 0.00 0.80
DSA23: MKC 0.00 0.00 0.00 0.13
$4$DUA0: (RABBIT) SYSTEM_0 0.20 0.65 0.17 1.97
$4$DUA2: (RABBIT) SYSTEM_0 0.20 0.65 0.17 1.97
$4$DUA3: (RABBIT) SYSTEM_1 0.00 0.14 0.00 2.49

PLAYBACK SUMMARIZING

```

18.8.7 Using Remote Playback Monitoring

If suitably privileged, you can collect MONITOR data from any system to which your system has a DECnet connection. You can then display the data live on your local system. To do so, follow these steps:

1. In the default DECnet directory on each remote system, create a file named MONITOR.COM, similar to the following:

```

$ !
$ !      * Enable MONITOR remote playback *
$ !
$ MONITOR /NODISPLAY/RECORD=SYS$NET ALL_CLASSES

```

2. On your local system, define a logical name for the remote system from which you want to collect data. Use the following syntax:

```
DEFINE remotenodename_mon node::task=monitor
```

You might want to define, in a login command procedure, a series of logical names for all the systems you want to access.

3. To display the remote MONITOR data as it is being collected, enter a command using the following syntax:

```
MONITOR/INPUT=remotenodename_mon classnames
```

You can also place MONITOR.COM files in directories other than the default DECnet directory and use access control strings or proxy accounts to invoke these command files remotely.

When you invoke MONITOR on your local system, a process is created on the remote system that executes the MONITOR.COM command file. The remote system therefore experiences some associated CPU and DECnet overhead. You can regulate the overhead in the MONITOR.COM file by using the /INTERVAL qualifier and the list of class names.

Getting Information About the System

18.8 Monitoring Operating System Performance

Section 18.8.10 describes remote monitoring in a mixed version cluster system.

18.8.8 Rerecording Monitoring

Rerecording is a combination of playback and recording. You can use it for data reduction of recording files. When you play back an existing recording file, all MONITOR options are available to you; thus, you can choose to record a subset of the recorded classes and a subset of the recorded time segment and a larger interval value.

All these techniques produce a new, smaller recording file at the expense of some of the recorded data. A larger interval value reduces the volume of the collected data, so displays and summary output produced from the newer recorded file will be less precise. Note that average rate values are not affected in this case, but average level values are less precise (since the sample size is reduced), as are maximum and minimum values. The following example shows how to use the rerecording mode of operation:

Example

```
$ SUBMIT MONREC.COM
```

MONREC.COM contains the following commands:

```
$ MONITOR/NODISPLAY/RECORD/INTERVAL=60 /BEGINNING=8:00/ENDING=16:00 DECNET,LOCK  
$ MONITOR/INPUT/NODISPLAY/RECORD DECNET
```

The first command runs in a batch job, recording DECnet and lock management data once every minute between the hours of 8 A.M. and 4 P.M.. The second command, which is issued after the first command completes, rerecords the data by creating a new version of the MONITOR.DAT file, containing only the DECnet data.

18.8.9 Running MONITOR Continuously

You can develop a database of performance information for your system by running MONITOR continuously as a background process. This section contains examples of procedures that you, as cluster manager, might use to create multifile clusterwide summaries.

You can adapt the command procedures to suit conditions at your site. Note that you must define the logical names SYS\$MONITOR and MON\$ARCHIVE in SYSTARTUP.COM before executing any of the command files.

The directory with the logical name SYS\$EXAMPLES includes three command procedures that you can use to establish the database. Instructions for installing and running the procedures are in the comments at the beginning of each procedure. Table 18-9 contains a brief summary of these procedures.

Table 18-9 MONITOR Command Procedures

Procedure	Description
MONITOR.COM	Creates a summary file from the recording file of the previous boot, and then begins recording for this boot. The recording interval is 10 minutes.

(continued on next page)

Getting Information About the System

18.8 Monitoring Operating System Performance

Table 18–9 (Cont.) MONITOR Command Procedures

Procedure	Description
MONSUM.COM	Generates two clusterwide multifile summary reports that are mailed to the system manager: one report is for the previous 24 hours, and the other is for the previous day's prime-time period (9 A.M. to 6 P.M.). The procedure resubmits itself to run each day at midnight.
SUBMON.COM	Starts MONITOR.COM as a detached process. Invoke SUBMON.COM from the site-specific startup command procedure.

While MONITOR records data continuously, a summary report can cover any finite time segment. The MONSUM.COM command procedure, which is executed every midnight, produces and mails the two multifile summary reports described in Table 18–9. Because these reports are not saved as files, to keep them, you must either extract them from your mail file or alter the MONSUM.COM command procedure to save them.

18.8.9.1 Using the MONITOR.COM Procedure

The procedure in Example 18–7 archives the recording file and summary file from the previous boot and initiates continuous recording for the current boot. (Note that this procedure does not purge recording files.)

Example 18–7 MONITOR.COM Procedure

```
$ SET VERIFY
$ !
$ ! MONITOR.COM
$ !
$ ! This command file is to be placed in a cluster-accessible directory
$ ! called SYS$MONITOR and submitted at system startup time as a detached
$ ! process via SUBMON.COM. For each node, MONITOR.COM creates, in
$ ! SYS$MONITOR, a MONITOR recording file that is updated throughout the
$ ! life of the boot. It also creates, in MON$ARCHIVE, a summary file from
$ ! the recording file of the previous boot, along with a copy of that
$ ! recording file. Include logical name definitions for both cluster-
$ ! accessible directories, SYS$MONITOR and MON$ARCHIVE, in SYSTARTUP.COM.
$ !
$ SET DEF SYS$MONITOR
$ SET NOON
$ PURGE MONITOR.LOG/KEEP:2
$ !
$ ! Compute executing node name and recording and summary file names
$ ! (incorporating node name and date).
$ !
$ NODE = F$GETSYI("NODENAME")
$ SEP = ""
$ IF NODE .NES. "" THEN SEP = "_"
$ DAY = F$EXTRACT(0,2,F$TIME())
$ IF F$EXTRACT(0,1,DAY) .EQS. " " THEN DAY = F$EXTRACT(1,1,DAY)
$ MONTH = F$EXTRACT(3,3,F$TIME())
```

(continued on next page)

Getting Information About the System

18.8 Monitoring Operating System Performance

Example 18-7 (Cont.) MONITOR.COM Procedure

```
$ ARCHFILNAM = "MON$ARCHIVE:" + NODE + SEP + "MON" + DAY + MONTH
$ RECFIL = NODE + SEP + "MON.DAT"
$ SUMFIL = ARCHFILNAM + ".SUM"

$ !
$ ! Check for existence of recording file from previous boot and skip
$ ! summary if not present.
$ !
$ OPEN/READ/ERROR=NORECFIL RECORDING 'RECFIL'
$ CLOSE RECORDING
$ !
$ !
$ ! Generate summary file from previous boot.
$ !
$ MONITOR /INPUT='RECFIL' /NODISPLAY /SUMMARY='SUMFIL' -
$ ALL_CLASSES+MODE/ALL+STATES/ALL+SCS/ITEM=ALL+SYSTEM/ALL+DISK/ITEM=ALL
$ !
$ !
$ ! Compute subject string and mail summary file to cluster manager.
$ !
$ !
$ A=""
$ B=" MONITOR Summary "
$ SUB = A + NODE + B + F$TIME() + A
$ MAIL/SUBJECT='SUB' 'SUMFIL' CLUSTER_MANAGER
$ !
$ !
$ ! Archive recording file and delete it from SYS$MONITOR.
$ !
$ COPY 'RECFIL' 'ARCHFILNAM'.DAT
$ DELETE 'RECFIL';*
$ !
$ NORECFIL:
$ SET PROCESS/PRIORITY=15
$ !
$ !
$ ! Begin recording for this boot. The specified /INTERVAL value is
$ ! adequate for long-term summaries; you might need a smaller value
$ ! to get reasonable "semi-live" playback summaries (at the expense
$ ! of more disk space for the recording file).
$ !
$ MONITOR /INTERVAL=300 /NODISPLAY /RECORD='RECFIL' ALL_CLASSES
$ !
$ !
$ ! End of MONITOR.COM
$ !
```

18.8.9.2 Using the SUBMON.COM Procedure

The procedure in Example 18-8 submits MONITOR.COM as a detached process from SYSTARTUP.COM to initiate continuous recording for the current boot.

Example 18-8 SUBMON.COM Procedure

```
$ SET VERIFY
$ !
$ ! SUBMON.COM
```

(continued on next page)

Getting Information About the System

18.8 Monitoring Operating System Performance

Example 18-8 (Cont.) SUBMON.COM Procedure

```
$ !
$ ! This command file is to be placed in a cluster-accessible directory
$ ! called SYS$MONITOR. At system startup time, for each node, it is
$ ! executed by SYSTARTUP.COM, following logical name definitions for
$ ! the cluster-accessible directories SYS$MONITOR and MON$ARCHIVE.
$ !
$ !
$ ! Submit detached MONITOR process to do continuous recording.
$ !
$ !
$ RUN  SYS$SYSTEM:LOGINOUT.EXE -
        /UIC=[1,4] -
        /INPUT=SYS$MONITOR:MONITOR.COM -
        /OUTPUT=SYS$MONITOR:MONITOR.LOG -
        /ERROR=SYS$MONITOR:MONITOR.LOG -
        /PROCESS_NAME="Monitor" -
        /WORKING_SET=512 -
        /MAXIMUM_WORKING_SET=512 -
        /EXTENT=512/NOSWAPPING
$ !
$ !
$ ! End of SUBMON.COM
$ !
```

18.8.9.3 Using the MONSUM.COM Procedure

The procedure in Example 18-9 produces daily and prime-time clusterwide summaries.

Example 18-9 MONSUM.COM Procedure

```
$ SET VERIFY
$ !
$ ! MONSUM.COM
$ !
$ ! This command file is to be placed in a cluster-accessible directory
$ ! called SYS$MONITOR and executed at the convenience of the cluster
$ ! manager. The file generates both 24-hour and "prime time" cluster
$ ! summaries and resubmits itself to run each day at midnight.
$ !
$ SET DEF SYS$MONITOR
$ SET NOON
$ !
$ ! Compute file specification for MONSUM.COM and resubmit the file.
$ !
$ FILE = F$ENVIRONMENT("PROCEDURE")
$ FILE = F$PARSE(FILE,,, "DEVICE")+F$PARSE(FILE,,, "DIRECTORY")+F$PARSE(FILE,,, "NAME")
$ SUBMIT 'FILE' /AFTER=TOMORROW /NOPRINT
$ !
$ ! Generate 24-hour cluster summary.
$ !
$ !
$ MONITOR/INPUT=(SYS$MONITOR:*MON*.DAT;* , MON$ARCHIVE:*MON*.DAT;*) -
  /NODISPLAY/SUMMARY=MONSUM.SUM -
  ALL_CLASSES+DISK/ITEM=ALL+SCS/ITEM=ALL-
  /BEGIN="YESTERDAY+0:0:0.00" /END="TODAY+0:0:0.00" /BY_NODE
```

(continued on next page)

Getting Information About the System

18.8 Monitoring Operating System Performance

Example 18-9 (Cont.) MONSUM.COM Procedure

```
$ !
$ !
$ ! Mail 24-hour summary file to cluster manager and delete the file from
$ ! SYS$MONITOR.
$ !
$ !
$ MAIL/SUBJECT="Daily Monitor Clusterwide Summary" MONSUM.SUM CLUSTER_MANAGER
$ DELETE MONSUM.SUM;*
$ !
$ ! Generate prime-time cluster summary.
$ !
$ !
$ MONITOR/INPUT=(SYS$MONITOR:*MON*.DAT;* ,MON$ARCHIVE:*MON*.DAT;*) -
$ /NODISPLAY/SUMMARY=MONSUM.SUM -
$ ALL_CLASSES+DISK/ITEM=ALL+SCS/ITEM=ALL-
$ /BEGIN="YESTERDAY+9:0:0.00" /END="YESTERDAY+18:0:0.00" /BY_NODE
$ !
$ !
$ ! Mail prime-time summary file to cluster manager and delete the file
$ ! from SYS$MONITOR.
$ !
$ !
$ MAIL/SUBJECT="Prime-Time Monitor Clusterwide Summary" MONSUM.SUM CLUSTER_MANAGER
$ DELETE MONSUM.SUM;*
$ !
$ ! End of MONSUM.COM
$ !
```

Note that MAIL commands in this procedure send files to user CLUSTER_MANAGER. Replace CLUSTER_MANAGER with the appropriate user name or logical name for your site.

Because summary data might be extensive, Digital recommends that you print out summary files.

18.8.10 Remote Monitoring in a Mixed-Version VMSccluster System

Remote monitoring is a feature of the Monitor utility (MONITOR) that enables you to monitor any node in a VMSccluster system. You can do this either by issuing the MONITOR CLUSTER command or by adding the /NODE qualifier to any interactive MONITOR request.

Remote monitoring in a VMSccluster system may not be compatible between nodes that are running different OpenVMS versions. Table 18-10 shows the compatibility of versions for remote monitoring.

Table 18-10 Remote Monitoring Compatibility in a VMSccluster System

	OpenVMS Alpha and VAX Version 6.n	OpenVMS Alpha Version 1.5 and VAX Version 5.n
OpenVMS Alpha and VAX Version 6.n	Yes	No
OpenVMS Alpha Version 1.5 and VAX Version 5.n	No	Yes

Getting Information About the System

18.8 Monitoring Operating System Performance

If you attempt to monitor a remote node that is incompatible, the following message is displayed:

```
%MONITOR-E-SRVMISMATCH, MONITOR server on remote node is an incompatible version
```

If you receive this message, you can use MONITOR to obtain data about the remote node. To do this, record the data on the remote node and then run the MONITOR playback feature to examine the data on the local node.

Another difference exists when you are monitoring remote nodes in a VMScLuster system. The limit on the number of disks that can be monitored on OpenVMS VAX Version 6.2 and OpenVMS Alpha Version 6.2 was raised from 799 to 909 for record output and from 799 to 1817 for display and summary outputs. If you are monitoring a remote node running OpenVMS Version 6.2 from a system running an earlier version of OpenVMS, its limit of 799 is in effect.

For more information on MONITOR, see the *OpenVMS System Management Utilities Reference Manual*.

Tracking Resource Use

This chapter describes how to find out how your system resources have been used. You can use this information to:

- Charge users for the resources they have used. You can produce reports of the resources used by individual users.
- Plan your future equipment requirements. You can monitor changing patterns of resource use and predict future demands.
- Troubleshoot the system. You can check the final exit status of processes.
- Improve system performance. You can find out the load that individual images and processes place on your system.
- Detect security breaches. You can identify unusual patterns of resource use.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Determining which resources are being tracked	Section 19.2
Controlling which resources are tracked	Section 19.3
Starting up a new accounting file	Section 19.4
Moving the accounting file	Section 19.5
Producing reports of resource use	Section 19.6
Setting up accounting groups	Section 19.7
Monitoring disk space	Section 19.8

This chapter explains the following concept:

Concept	Section
Accounting files	Section 19.1

19.1 Understanding Accounting Files

The system gathers information on resource use. For example, the information can include the resources such as CPU time used by each print job. The system stores this information in **accounting files**.

The resources tracked by default depend on the model of computer you use. However, you can control which resources are tracked. If you do not want to track resource use, you can stop the accounting file tracking resource use altogether. (See Section 19.3.)

Tracking Resource Use

19.1 Understanding Accounting Files

Each node in a VMScluster has its own accounting file, known as its **current accounting file**. By default, this file is SYS\$MANAGER:ACCOUNTNG.DAT, but you can control which file is used (see Section 19.5).

The information in the accounting files is in binary. You cannot display it with the TYPE command. To display the information, you use the Accounting utility (ACCOUNTING). (See Section 19.6.)

19.2 Determining Which Resources Are Being Tracked

To determine which resources are currently being tracked, use the SHOW ACCOUNTING command:

```
$ SHOW ACCOUNTING
```

This command produces a screen display (see the example) that contains keywords in the following two categories:

- Keywords that show which types of resource are being tracked:

Keyword	Type of Resource
IMAGE	Resources used by an image
LOGIN_FAILURE	Resources used by an unsuccessful attempt to log in
MESSAGE	Unformatted resource record written to the accounting file by a call to the \$SNDJBC system service
PRINT	Resources used by a print job
PROCESS	Resources used by a process

- Keywords that show which types of process are being tracked. When the resources for processes or images are tracked, these keywords show the process type:

Keyword	Type of Process
BATCH	Batch process
DETACHED	Detached process
INTERACTIVE	Interactive process
NETWORK	Network process
SUBPROCESS	Subprocess (the parent process can be a batch, detached, interactive, or network process)

Example

```
$ SHOW ACCOUNTING
```

Accounting is currently enabled to log the following activities:

```
PROCESS      any process termination
IMAGE        image execution
INTERACTIVE  interactive job termination
LOGIN_FAILURE login failures
NETWORK      network job termination
PRINT        all print jobs
```

The keywords in this example show that the local node is tracking the resources used by each:

- Interactive and network process

Tracking Resource Use

19.2 Determining Which Resources Are Being Tracked

- Image running in an interactive or network process
- Login failure
- Print job

19.3 Controlling Which Resources Are Tracked

You can control which resources the system tracks. To save disk space, you can stop the system tracking resources you are not interested in.

How to Perform This Task

1. Use the SET ACCOUNTING command with the /ENABLE and /DISABLE qualifiers in the following format to control which resources are tracked:

```
SET ACCOUNTING/DISABLE[=(keyword[,...])]/ENABLE[=(keyword[,...])]
```

The keywords are the same as those explained in Section 19.2.

2. If you want to make this change permanent, edit the SET ACCOUNTING command in the SYS\$MANAGER:SYSTART_VMS.COM startup file.

Example

This example prevents the tracking of all resources except those used by interactive and batch processes:

```
$ SET ACCOUNTING/DISABLE/ENABLE= (PROCESS, INTERACTIVE, BATCH)
```

The /DISABLE qualifier is not followed by a keyword. Therefore, the qualifier disables the tracking of all resources. The /ENABLE qualifier then enables the tracking of the resources used by interactive and batch processes.

19.4 Starting Up a New Accounting File

To start up a new current accounting file, use the following command:

```
$ SET ACCOUNTING/NEW_FILE
```

This closes the current accounting file and opens a new version of it.

If the system encounters an error when trying to write information to the current accounting file, it automatically closes the file and opens a new version of it.

Example

This example closes the current accounting file, opens a new version of it, and changes the name of the old file to WEEK_24_RESOURCES.DAT. You can retain this file as a record of the resources used in that week.

```
$ SET ACCOUNTING/NEW_FILE  
$ RENAME SYS$MANAGER:ACCOUNTNG.DAT;-1 WEEK_24_RESOURCES.DAT
```

19.5 Moving the Accounting File

When you first install your system, the current accounting file is SYS\$MANAGER:ACCOUNTNG.DAT.

This file can become quite large. Moving it from your system disk can improve system performance.

Tracking Resource Use

19.5 Moving the Accounting File

How to Perform This Task

1. Define the logical name ACCOUNTNG in your system logical name table to point to the file you want to use. For example:

```
$ DEFINE ACCOUNTNG MYDISK:[MYDIR]MYFILE.DAT/SYSTEM
```

Give the full file specification, including the device and directory.

Note

Two nodes cannot log information in the same accounting file. If you define ACCOUNTNG on two nodes to point to the same file, each node will open and use its own version of the file.

2. To make the change permanent, add this definition to the file SYS\$MANAGER:SYLOGICALS.COM.
3. Use the SET ACCOUNTING command with the /NEW_FILE qualifier to create and use the new file:

```
$ SET ACCOUNTING/NEW_FILE
```

Example

This example changes the current accounting file to [MYDIR]MYDISK:MYFILE.DAT.

```
$ DEFINE ACCOUNTNG MYDISK:[MYDIR]MYFILE.DAT/SYSTEM
$ SET ACCOUNTING/NEW_FILE
```

19.6 Producing Reports of Resource Use

The three types of reports are:

Type of Report	Qualifier
Brief	/BRIEF (the default)
Full	/FULL
Summary	/SUMMARY

To produce a report, use the ACCOUNTING command with the appropriate qualifier in the following format:

```
ACCOUNTING [filespec[...]/qualifier[...]]
```

This runs the Accounting utility. The *filespec* parameter lists the accounting files you want to process. If you omit it, the Accounting utility processes the default current accounting file, SYS\$MANAGER:ACCOUNTNG.DAT.

By default, the Accounting utility processes all the records in the accounting files you specify. You can use selection qualifiers to specify which records you want to process.

By default, brief and full reports present the records in the order in which they were logged in the accounting file. When you produce brief and full reports, you can use the /SORT qualifier to specify another order.

Tracking Resource Use

19.6 Producing Reports of Resource Use

Example

This example produces a brief report of the information in the file that the logical name ACCOUNTNG points to. The /TYPE qualifier selects records for print jobs only. The /SORT qualifier displays them in reverse alphabetical order of user name.

```
$ ACCOUNTING ACCOUNTNG/TYPE=PRINT/SORT=-USER
```

Date / Time	Type	Subtype	Username	ID	Source	Status
13-APR-1995 13:36:04	PRINT		SYSTEM	20A00442		00000001
13-APR-1995 12:42:37	PRINT		JONES	20A00443		00000001
13-APR-1995 14:43:56	PRINT		FISH	20A00456		00000001
14-APR-1995 19:39:01	PRINT		FISH	20A00265		00000001
14-APR-1995 20:09:03	PRINT		EDWARDS	20A00127		00000001
14-APR-1995 20:34:45	PRINT		DARNELL	20A00121		00000001
14-APR-1995 11:23:34	PRINT		CLARK	20A0032E		00040001
14-APR-1995 16:43:16	PRINT		BIRD	20A00070		00040001
14-APR-1995 09:30:21	PRINT		ANDERS	20A00530		00040001

19.7 Setting Up Accounting Groups

Users are already organized into UIC security groups. For accounting purposes, security groups are often inappropriate. You can put users into accounting groups with the Authorize utility using the /ACCOUNT qualifier. In this way, each user is in an accounting group and a security group.

Using the Accounting utility, you can:

- Summarize the resources used by all the users in a particular accounting or security group. To do this, use the ACCOUNT or UIC keyword with the /SUMMARY qualifier.
- Select records for all the users in a particular accounting or security group. To do this, use the /ACCOUNT or /UIC qualifier.

How to Perform This Task

1. Plan your accounting groups. Decide which users you want in each accounting group, and choose names for the groups.

The name of an accounting group can be a maximum of eight characters long.

2. Change the account field values in the UAF. Use the Authorize utility's MODIFY command in the following format to change the value in the account field to the name of the user's accounting group:

```
MODIFY username/ACCOUNT=accounting-group-name
```

where:

username is the name of the user

accounting-group-name is the name of the accounting group that you want that user to be in

The next time your users log in, they will be in their new accounting groups, and their resource use will be tagged with the appropriate accounting group names.

Tracking Resource Use

19.7 Setting Up Accounting Groups

Example

This example modifies the accounting group name to SALES_W8 for the username FORD:

```
$ RUN SYS$SYSTEM:AUTHORIZE
UAF> MODIFY FORD/ACCOUNT=SALES_W8
UAF> EXIT
```

19.8 Monitoring Disk Space

To find out how much disk space a user is using, use SYSMAN or, if you have not enabled disk quotas, the DIRECTORY command.

How to Perform This Task

Use either of the following methods:

- Use the SYSMAN command DISKQUOTA SHOW in the following format:

```
DISKQUOTA SHOW uic [/DEVICE=diskname]
```

This shows the number of blocks used by all the files that are owned by the specified user on the specified disk.

- Use the DIRECTORY command with the /SIZE and /GRAND_TOTAL qualifiers in the following format:

```
DIRECTORY diskname:[username...]/SIZE=ALLOCATION/GRAND_TOTAL
```

This shows the number of blocks used by all the files in and under the specified user's root directory.

Note that the DIRECTORY command does not include the blocks used by file headers or the user's root directory.

Examples

1. This example uses SYSMAN to find out the number of blocks used by all the files that are owned by each user.

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> DISKQUOTA SHOW *

%SYSMAN-I-QUOTA, disk quota statistics on device SYS$SYSTEM:MYDISK
Node UNION
  UIC          Usage      Permanent Quota  Overdraft Limit
[0,0]          0          1000             100
[DOC,EDWARDS]  115354     150000           5000
[DOC,FISH]     177988     250000           5000
[DOC,SMITH]    140051     175000           5000
[DOC,JONES]    263056     300000           5000
```

2. This example uses the DIRECTORY command to show the number of blocks allocated by all the files in and under MYDISK:[PARSONS].

```
$ DIRECTORY MYDISK:[PARSONS...]/SIZE=ALLOCATION/GRAND_TOTAL
Grand total of 28 directories, 2546 files, 113565 blocks.
```


VMSccluster Considerations

This chapter describes concepts related to the VMSccluster environment; it also tells how the Show Cluster utility (SHOW CLUSTER) can display information about a cluster and how the System Management utility (SYSMAN) can help you manage a VMSccluster environment.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Beginning to use SHOW CLUSTER commands	Section 20.3.2
Adding information to a report	Section 20.3.3
Controlling the display of data	Section 20.3.4
Formatting the display of data	Section 20.3.5
Creating a startup initialization file	Section 20.3.6
Using command procedures containing SHOW CLUSTER commands	Section 20.3.7
Using SYSMAN to manage security and system time	Section 20.5
Using the SYSMAN DO command to manage a VMSccluster	Section 20.6

This chapter explains the following concepts:

Concept	Section
VMSccluster systems	Section 20.1
Setting up a VMSccluster environment	Section 20.1.1
Clusterwide system management	Section 20.1.2
The Show Cluster utility (SHOW CLUSTER)	Section 20.3.1
SYSMAN and VMSccluster management	Section 20.4

20.1 Understanding VMSccluster Systems

A **VMSccluster** system is a loosely coupled configuration of two or more computers and storage subsystems. A VMSccluster system appears as a single system to the user even though it shares some or all of the system resources. When a group of computers shares resources clusterwide, the storage and computing resources of all of the computers are combined, which can increase the processing capability, communications, and availability of your computing system.

A **shared resource** is a resource (such as a disk) that can be accessed and used by any node in a VMSccluster system. Data files, application programs, and printers are just a few items that can be accessed by users on a cluster with

VMScLuster Considerations

20.1 Understanding VMScLuster Systems

shared resources, without regard to the particular node on which the files or program or printer might physically reside.

When disks are set up as shared resources in a VMScLuster environment, users have the same environment (password, privileges, access to default login disks, and so on) regardless of the node that is used for logging in. You can realize a more efficient use of mass storage with shared disks, because the information on any device can be used by more than one node—the information does not have to be rewritten in many places. You can use the OpenVMS MSCP, which is the mass storage control protocol, or TMSCP, which is the tape mass storage control protocol, server software to make tapes accessible to nodes that are not directly connected to the storage devices.

You can also set up print and batch queues as shared resources. In a VMScLuster configuration with shared print and batch queues, a single queue database manages the queues for all nodes. The queue database makes the queues available from any node. For example, suppose your VMScLuster configuration has fully shared resources and includes nodes ALBANY, BASEL, and CAIRO. A user logged in to node ALBANY can send a file that physically resides on node BASEL to a printer that is physically connected to node CAIRO, and the user never has to specify (or even know) the nodes for either the file or the printer.

A number of types of VMScLuster configurations are possible. Refer to *Guidelines for VMScLuster Configurations* and either the VMScLuster or VAXcluster Software Product Description (SPD) for complete information about supported devices and configurations.

The following sections briefly describe VMScLuster systems. For complete information about setting up and using a VMScLuster environment, see *VMScLuster Systems for OpenVMS*.

20.1.1 Setting Up a VMScLuster Environment

Once you have planned your configuration, installed the necessary hardware, and checked hardware devices for proper operation, you can set up a VMScLuster system using various system software facilities. Setup procedures to build your VMScLuster system follow.

Procedure	For More Information
Installing or upgrading the operating system on the first VMScLuster computer	Installation and operations guide for your computer
Installing required software licenses	<i>OpenVMS License Management Utility Manual</i>
Configuring and starting the DECnet for OpenVMS network	<i>DECnet for OpenVMS Networking Manual</i>
Preparing files that define the cluster operating environment and that control disk and queue operations	<i>VMScLuster Systems for OpenVMS</i>
Adding computers to the cluster	<i>VMScLuster Systems for OpenVMS</i>

Depending on various factors, the order in which these operations are performed can vary from site to site, as well as from cluster to cluster at the same site.

20.1.2 Clusterwide System Management

Once any system is installed, the system manager must decide how to manage users and resources for maximum productivity and efficiency while maintaining the necessary security. VMScLuster systems provide the flexibility to distribute users and resources to suit the needs of the environment. VMScLuster system resources can also be easily redistributed as needs change. Even with the vast number of resources available, the VMScLuster configuration can be managed as a single system.

VMScLuster system managers have several tools and products to help them manage their systems as a unified entity.

VMScLuster Tools

The following utilities are provided with the operating system:

Utility	Description
DECamds	Collects and analyzes data from multiple nodes simultaneously, directing all output to a centralized DECwindows display. (Refer to Section 20.2 and the <i>DECamds User's Guide</i> .)
Monitor utility	Provides basic performance data. (See Section 18.8.)
Show Cluster utility (SHOW CLUSTER)	Monitors activity in a VMScLuster configuration, and then collects and sends information about that activity to a terminal or other output device. (Described in Section 20.3.)
System Management utility (SYSMAN)	Allows the system manager to send common control commands across all, or a subset of, the nodes in the VMScLuster system. (Described in Section 20.6.)

System Management Applications

The following products are *not* provided with the operating system:

Product	Description
POLYCENTER solutions	A comprehensive set of operations management products and services to help you manage complex distributed environments. However, the POLYCENTER Software Installation utility is described in this manual in Section 3.7.
†Storage Library System (SLS) for VAX ‡Storage Library System (SLS) for Alpha	A set of software tools that enables tape, cartridge tape, and optical disks.
VMScLuster Console System (VCS)	Designed to consolidate the console management of the VMScLuster system at a single console terminal.
†VAX specific ‡Alpha specific	

You can find additional information about these system management tools in the appropriate product documentation.

20.2 Using DECamds to Analyze Data

The Digital Availability Manager for Distributed Systems (DECamds) is a real-time monitoring, diagnostic, and correction tool that assists system managers to improve OpenVMS system and VMScLuster availability. DECamds can help system programmers and analysts target a specific node or process for detailed analysis, and can help system operators and service technicians resolve hardware and software problems.

VMScluster Considerations

20.2 Using DECamsds to Analyze Data

DECamsds simultaneously collects and analyzes system data and process data from multiple nodes and displays the output on a DECwindows Motif display. Based on the collected data, DECamsds analyzes, detects, and proposes actions to correct resource and denial issues in real-time.

For more information, see the *DECamsds User's Guide*.

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

The Show Cluster utility (SHOW CLUSTER) monitors nodes in a VMScluster system. You can use the utility to display information about cluster activity and performance.

The following sections describe the Show Cluster utility and explain how to perform these tasks:

Task	Section
Begin to use SHOW CLUSTER commands	Section 20.3.2
Add information to a report	Section 20.3.3
Control the display of data	Section 20.3.4
Create a startup initialization file	Section 20.3.6
Use command procedures containing SHOW CLUSTER commands	Section 20.3.7

20.3.1 Understanding the Show Cluster Utility

You can display SHOW CLUSTER information on your terminal screen or send it to a device or a file. You can use the Show Cluster utility interactively, with command procedures, or with an initialization file in which you define default settings. Because this utility is installed with the CMKRNL privilege, SHOW CLUSTER requires no special privilege.

SHOW CLUSTER information includes approximately 100 fields of data. You can customize the appearance of SHOW CLUSTER reports or define reports for access to often-needed data.

SHOW CLUSTER reports are organized by classes and fields:

Unit of Organization	Description
Class	Group of related fields of information. You can use class names to selectively add or remove an entire class from a report. Each class displays certain fields by default. Some classes have additional fields that you can add or remove using the field name.
Field	Column of data in a report. You use a unique field name to refer to each field of data. You can use the field name to selectively add or remove a field from reports. For the names and descriptions of all of the fields in each class, see the ADD (Field) command in the <i>OpenVMS System Management Utilities Reference Manual</i> .

You can add fields or classes to the default SHOW CLUSTER report. If you add a field or class to a report in a continuous display, SHOW CLUSTER automatically adds the new data to the display.

VMSCluster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

Example 20-1 shows a sample default SHOW CLUSTER report. The default report has two classes of information: SYSTEMS and MEMBERS. Below each class name are columns of fields that are associated with each class of information.

Example 20-1 SHOW CLUSTER Default Display

View of Cluster from system ID 77777 node: CLUB

SYSTEMS			MEMBERS
NODE	HW_TYPE	SOFTWARE	STATUS
CLUB	DEC 4000 Model 610	VMS X5EM	MEMBER
DISK12	RF72	RFX T251	
CONS07	EVAX	CON V1.0	
DISK14	RF72	RFX V255	
CHIP	DEC 4000 Model 620	VMS X5EM	MEMBER
DISK3	RF72	RFX V254	
DISK1	RF72	RFX V256	
SPREE	DEC 3000 Model 500	VMS X5EM	MEMBER
SPRITZ	VAX 4000-300	VMS A5.5	MEMBER

Table 20-1 briefly describes the fields shown in Example 20-1.

Table 20-1 Fields in Default SHOW CLUSTER Report

Field Name	Description
NODE	Node name of the remote system. Normally, the cluster manager sets the node name using the system parameter SCSNODE. The node name should be the same as the DECnet for OpenVMS node name.
HW_TYPE	Hardware type and model of the remote system.
SOFTWARE	Name and version of the operating system currently running on the remote system.
STATUS	Status of the node in the cluster. (MEMBER indicates that the system is participating in the cluster.)

Over time, you can determine the most valuable classes and fields of data for your SHOW CLUSTER reports; you can then create a startup initialization file that establishes your default report formats. You can also build command procedures to use while running SHOW CLUSTER interactively. In this way, you can quickly reformat the report to show the data that is relevant for your installation. Startup initialization files and command procedures are explained later in this chapter.

Because SHOW CLUSTER information includes many fields of data, the report can quickly extend beyond screen limits. Therefore, SHOW CLUSTER provides mechanisms to help you control the display of data, including the following:

- 38 SHOW CLUSTER commands
- A default keypad, which you can redefine

These mechanisms are described in detail in the *OpenVMS System Management Utilities Reference Manual*.

VMSccluster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

20.3.2 Beginning to Use SHOW CLUSTER Commands

To use the Show Cluster utility, you enter the SHOW CLUSTER command. If you specify the command without any qualifiers, however, SHOW CLUSTER simply displays a default report like that shown in Example 20–1 and then displays the DCL prompt.

In a continuous display, on the other hand, you can enter SHOW CLUSTER commands to control report output. You can, for example, add classes or fields to, or remove classes or fields from, reports. To invoke a continuous display, in which you can enter SHOW CLUSTER commands, use the /CONTINUOUS qualifier on the SHOW CLUSTER command. (SHOW CLUSTER command qualifiers are described in Section 20.3.2.3.)

How to Perform This Task

To invoke a continuous display of default SHOW CLUSTER report information, enter the following command:

```
$ SHOW CLUSTER/CONTINUOUS
```

SHOW CLUSTER then displays a default report. By default, SHOW CLUSTER updates the display every 15 seconds, with the changed data displayed in reverse video. After the default report, SHOW CLUSTER displays the following prompt:

Command>

(If the report extends below the limit of your terminal screen and you do not see the Command> prompt, you can press Return to display the prompt.)

The following sections contain instructions for performing beginning SHOW CLUSTER tasks:

Task	Section
Viewing information that is off the screen	Section 20.3.2.1
Exiting from a continuous display	Section 20.3.2.2
Using SHOW CLUSTER qualifiers	Section 20.3.2.3

20.3.2.1 Viewing Information That Is Off the Screen

The PAN command allows you to view the entire display by shifting your view of the display by column (horizontally) or by line (vertically).

Note

Report headings also move out of view as the reports in the display are panned beyond the limits of the screen. The SCROLL command, which is explained in Section 20.3.5.4, preserves the headings as you scroll information. To use the SCROLL command, you must take the additional step of selecting a report if you have more than one report on the screen.

How to Perform This Task

To pan the display, do one of the following:

- Enter PAN commands at the command prompt; for example:

```
Command> PAN DOWN 10
```

The command in this example moves the display down 10 lines.

VMScLuster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

- Define the arrow keys as PAN commands:

Command> SET FUNCTION PAN

This command redefines the arrow keys as follows:

Arrow Key	Redefinition
↑	PAN UP 1
↓	PAN DOWN 1
→	PAN RIGHT 1
←	PAN LEFT 1

You can then use the arrow keys to move up, down, right, and left in the display.

See the SET FUNCTION and PAN commands in the *OpenVMS System Management Utilities Reference Manual* for details.

Resetting Arrow Keys

By default, the SHOW CLUSTER arrow keys are set to the EDIT function. This means that, at the command prompt, you can perform command line editing that is similar to DCL line-mode editing. For example, the left arrow key moves the cursor to the left, and the up arrow key recalls the previous command. See the *OpenVMS User's Manual* for information on DCL line-mode editing.

When you use the SET FUNCTION command, you reset the function keys. After that, the arrow keys are redefined and DCL line-mode editing is disabled.

To reset the arrow keys, enter the following command:

Command> SET FUNCTION EDIT

20.3.2.2 Exiting from a Continuous Display

To exit from a continuous display, do one of the following:

- To return to the DCL prompt, do one of the following:
 - Enter EXIT after the Command> prompt.
 - Press Ctrl/Z.
 - Press Ctrl/Y.
- To exit without erasing the screen, press Ctrl/C.

20.3.2.3 Using SHOW CLUSTER Qualifiers

Table 20–2 briefly describes the qualifiers you can use with the SHOW CLUSTER command. The *OpenVMS System Management Utilities Reference Manual* contains reference information about these SHOW CLUSTER qualifiers.

VMScluster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

Table 20-2 SHOW CLUSTER Qualifiers

Qualifier	Function
/BEGINNING= <i>time</i>	Specifies the time that the SHOW CLUSTER session is to begin.
/CONTINUOUS	Controls whether SHOW CLUSTER runs as a continuously updating display.
/ENDING= <i>time</i>	Specifies the time that the SHOW CLUSTER session is to end.
/INTERVAL= <i>seconds</i>	Specifies the number of seconds that report information remains on the screen before it is updated.
/OUTPUT= <i>file-spec</i>	Directs the output from SHOW CLUSTER to the specified file instead of to the current SYS\$OUTPUT device.

Example

In a continuous display, SHOW CLUSTER updates the display every 15 seconds by default. You can change this interval by using the /INTERVAL qualifier.

```
$ SHOW CLUSTER/CONTINUOUS/INTERVAL=5
```

In this example, SHOW CLUSTER updates reports every 5 seconds, displaying changed data in reverse video.

20.3.3 Adding Information to a Report

When you use the SHOW CLUSTER command, the resulting report is only part of the total information available. As shown in Example 20-1, the default classes displayed are MEMBERS and SYSTEMS. Table 20-3 briefly describes all the classes you can display in SHOW CLUSTER reports. See the *OpenVMS System Management Utilities Reference Manual* for details about these classes.

Table 20-3 Classes of Information Available in SHOW CLUSTER Reports

Classes	Information Displayed
CIRCUITS	Describes virtual circuits on VMScluster systems.
CLUSTER	Shows general information about the VMScluster system, such as the time it was formed, the last time a system joined or left, and the VMScluster quorum.
CONNECTIONS	Describes the connections established over a virtual circuit in the VMScluster system
COUNTERS	Shows counts of the total accumulated traffic over a connection for the life of the connection.
CREDITS	Shows send and receive credit counts for connections in the VMScluster system.
ERRORS	Displays a count of the errors on each port, along with information on the feasibility of reinitializing a port.
LOCAL_PORTS	Displays information on the local system interface to the VMScluster system, such as the name, number, and status of each port, and the number of entries in the queues associated with each port.
MEMBERS	Describes systems actively participating in the VMScluster system.

(continued on next page)

VMScLuster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

Table 20-3 (Cont.) Classes of Information Available in SHOW CLUSTER Reports

Classes	Information Displayed
SYSTEMS	Describes all VMScLuster systems. It shows node name, identification number, hardware type, and software version.

Example

The following example shows how to add the CLUSTER class to a SHOW CLUSTER display:

Command> ADD CLUSTER

Example 20-2 shows the display that results from entering the ADD CLUSTER command.

Example 20-2 SHOW CLUSTER Display with CLUSTER Report

View of Cluster from system ID 77777 node: CLUB

SYSTEMS			MEMBERS
NODE	HW_TYPE	SOFTWARE	STATUS
CLUB	DEC 4000 Model 610	VMS X5EM	MEMBER
DISK12	RF72	RFX T251	
CONS07	EVAX	CON V1.0	
DISK14	RF72	RFX V255	
CHIP	DEC 4000 Model 620	VMS X5EM	MEMBER
DISK3	RF72	RFX V254	
DISK1	RF72	RFX V256	
SPREE	DEC 3000 Model 500	VMS X5EM	MEMBER
SPRITZ	VAX 4000-300	VMS A5.5	MEMBER

CLUSTER		
CL_QUORUM	CL_VOTES	QD_NAME
2	3	

Table 20-1 describes the fields shown in the top section of the report shown in Example 20-2. Table 20-4 briefly describes the fields in the CLUSTER report.

Table 20-4 Fields in Sample CLUSTER Report

Field Name	Description
CL_QUORUM (Cluster quorum)	The number of votes that must be present for the cluster to function and permit user activity. CL_QUORUM is equal to (CL_EXPECTED_VOTES + 2) divided by 2.
CL_VOTES (Cluster votes)	Total number of votes contributed by all members of the cluster at any point in time.
QD_NAME (Quorum disk name)	Full device name of the quorum disk.

For detailed descriptions of the fields in the CLUSTER class, see the *OpenVMS System Management Utilities Reference Manual*.

VMScluster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

20.3.4 Controlling the Display of Data

Using SHOW CLUSTER commands, you can remove fields or classes from a display, remove broadcast messages from the screen, and refresh the screen display at any time. The following sections explain how to perform these operations.

20.3.4.1 Entering Commands to Display Data

SHOW CLUSTER allows you to customize the display of data during a continuous session by entering various commands. The *OpenVMS System Management Utilities Reference Manual* describes SHOW CLUSTER commands in detail.

Updating of the continuous display stops as soon as you enter input from the terminal keyboard. When you press the Return key after entering a command, updating of the display resumes until you enter another command.

By default, updating takes place at 15-second intervals. If you do not enter a new command within 15 seconds, the command prompt disappears, and two more lines of data take its place.

20.3.4.2 Removing Broadcast Messages

When you receive a system broadcast message during a continuous SHOW CLUSTER session, the message appears at the bottom of the screen. A multiline message fills as many lines of the screen as it needs.

How to Perform This Task

The last broadcast message you receive remains on the screen until you acknowledge it by entering input from the terminal in one of the following ways:

- Press the Return key.
- Refresh the screen by pressing Ctrl/W.
- Enter a command.

If you receive more than one broadcast message, SHOW CLUSTER waits until the next update interval to display the next message.

SHOW CLUSTER also displays error messages at the bottom of the screen. For an explanation of the error messages, see the *OpenVMS System Messages and Recovery Procedures Reference Manual*.

20.3.4.3 Refreshing the Screen

Ordinarily, a continuous display is updated or refreshed according to the default or specified interval time. SHOW CLUSTER scans the software databases, extracts and stores data for each field, displays any new or changed data, and updates the time. On Digital and Digital-compatible terminals, reverse video highlights any changed data.

How to Perform This Task

You can refresh the screen at any time by one of the following methods:

- Modify the format of the display with the ADD, REMOVE, INITIALIZE, or SET command.
- Use the REFRESH command.
- Press Ctrl/W.

VMScLuster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

20.3.5 Formatting the Display of Data

Because SHOW CLUSTER allows you to include additional fields and classes, you can produce reports that overflow the physical limits of the terminal screen. However, you can use a number of methods to modify the display to meet your needs:

Formatting Method	For More Information
Remove data from reports	Section 20.3.5.1
Modify field and screen size	Section 20.3.5.2
Move a report	Section 20.3.5.3
Scroll a report	Section 20.3.5.4

20.3.5.1 Removing Information from a Report

You might want to remove certain fields or classes to reduce the width of a report to fit the limits of your screen. Also, certain fields or classes might not be important for your particular needs. You can also remove particular types of data to reduce the length of the report.

How to Perform This Task

You use the REMOVE command to remove entire fields and classes, or subsets of fields and classes. To remove subsets of data, use the appropriate qualifier with the REMOVE *class-name* command. See the REMOVE commands in the *OpenVMS System Management Utilities Reference Manual* for appropriate class names and qualifiers.

Examples

1. Command> REMOVE SOFTWARE

The command in this example removes the SOFTWARE field from the SHOW CLUSTER report shown in Example 20-1.

See the ADD (Field) command description in the *OpenVMS System Management Utilities Reference Manual* for a list of valid field names.

2. Command> REMOVE MEMBERS

The command in this example removes the MEMBERS class from the SHOW CLUSTER report shown in Example 20-1.

20.3.5.2 Modifying Field and Screen Size

To make a report fit the physical limits of the screen, you can change the width of certain fields in the report. For example, if SHOW CLUSTER provides a field width that can contain any possible value and the values your cluster generates do not require that much space, you can adjust the field width with the SET (Field) command.

SHOW CLUSTER also allows you to adjust the size of the terminal screen. If the terminal is Digital-compatible and supports a wide report, you can set the screen to a width of up to 511 columns by specifying an appropriate value to the SET SCREEN command.

VMScluster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

Examples

1. Command> SET TRANSITION_TYPE/WIDTH=10

The command in this example sets the width of the TRANSITION_TYPE field to 10, which removes the time of day from the field but leaves the date.

2. Command> SET SCREEN=132

The command in this example sets the screen width to 132.

Refer to the *OpenVMS System Management Utilities Reference Manual* for more details about using the SET (Field) and SET SCREEN commands.

20.3.5.3 Moving a Report

By default, SHOW CLUSTER operates with AUTO_POSITIONING ON. This means that the utility automatically arranges the reports to take best advantage of the available display space. However, you can position reports manually with the MOVE command, which implicitly sets AUTO_POSITIONING to OFF.

If you have multiple reports in your display, you must first select the report to be repositioned. You use the command SELECT *window-name* to specify the report name; for example:

- SCS (the default report, which usually includes fields in the SYSTEMS and MEMBERS classes)
- CLUSTER
- LOCAL_PORTS

Note

To select any report except the default SCS report, you must first add the class to the display if it is not already displayed; for example:

Command> ADD LOCAL_PORTS

As an alternative, you can repeatedly press the Select function key or the period key on the keypad to cycle from one report to the next. The selected report appears highlighted.

How to Perform This Task

To move a report, do either of the following:

- Enter MOVE commands at the command prompt.
- Use the arrow keys that you define as MOVE commands.

Command> SET FUNCTION MOVE

This command redefines the arrow keys as follows:

Arrow Key	Redefinition
↑	MOVE UP 1
↓	MOVE DOWN 1
→	MOVE RIGHT 1

VMScIuster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

Arrow Key	Redefinition
←	MOVE LEFT 1

When you enter a MOVE command, the display changes position by column (horizontally) or by line (vertically). For example, entering the command MOVE LEFT 5 moves the display 5 columns to the left. An empty frame appears around the new position of the report.

When you are satisfied with the position of the report, enter the DESELECT command, which moves the report to the new position. Entering another SELECT command before the previous MOVE operation has been deselected also moves the report to its new position.

Example

```
Command> SELECT CLUSTER
Command> MOVE RIGHT 10
Command> DESELECT
```

Following is an explanation of the commands in the example:

1. The SELECT command selects the CLUSTER report (which is then highlighted).
2. The MOVE command positions the report frame 10 spaces to the right.
3. The DESELECT command terminates the MOVE operation and displays the contents of the report.

For more information, see the SELECT, SET FUNCTION, and DESELECT commands in the *OpenVMS System Management Utilities Reference Manual*.

To reset the arrow keys, enter the following command:

```
Command> SET FUNCTION EDIT
```

20.3.5.4 Scrolling a Report

The SCROLL command provides a means of quickly scanning through a report without losing column headings. Scrolling scans a display by field (horizontally) and by line (vertically). The report headings remain stationary when you scroll vertically.

When the display has more than one report, you must first select a report by entering the SELECT command. The selected report is highlighted.

How to Perform This Task

To scroll a display, do either of the following:

- Enter SCROLL commands at the command prompt.
- Use the arrow keys that you define as SCROLL commands.

```
Command> SET FUNCTION SCROLL
```

This command redefines the arrow keys as follows:

Arrow Key	Redefinition
↑	SCROLL UP 1
↓	SCROLL DOWN 1

VMScluster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

Arrow Key	Redefinition
→	SCROLL RIGHT 1
←	SCROLL LEFT 1

Example

```
Command> SELECT SCS
Command> SET FUNCTION SCROLL
```

The commands in this example first select the SCS report (which is then highlighted), and then set the arrow keys to scroll functions. See the SET FUNCTION and SCROLL commands in the *OpenVMS System Management Utilities Reference Manual* for more information.

To reset the arrow keys, enter the following command:

```
Command> SET FUNCTION EDIT
```

20.3.6 Creating a Startup Initialization File

To customize the SHOW CLUSTER display, you can create a startup initialization file, which the utility executes when you enter it. SHOW CLUSTER takes the original default display, and adds or removes whatever classes or fields you specify. The resulting display becomes your default startup format. A startup initialization file resembles the following:

```
!
!Startup Initialization File
!
!
INITIALIZE
REMOVE MEMBERS
ADD RP_REVISION,RP_TYPE,SYS_ID
SET SCREEN=132
```

This startup procedure causes SHOW CLUSTER to delete the MEMBERS class information from the default display. The procedure also adds the RP_REVISION and RP_TYPE fields from the CIRCUITS class and the SYS_ID field from the SYSTEMS class. The last line of the procedure sets the screen size to 132 columns.

How to Perform This Task

To create an initialization file, follow these steps:

1. Define the logical name SHOW_CLUSTER\$INIT as *device:[directory]SHCINI* before invoking SHOW CLUSTER.

For a startup file to execute before the display begins, you must assign the logical name SHOW_CLUSTER\$INIT to the initialization file; for example:

```
DEFINE SHOW_CLUSTER$INIT DEVA:[JONES]SHCINI
```

When invoked, SHOW CLUSTER searches for the file defined by SHOW_CLUSTER\$INIT. In this example, SHOW CLUSTER looks for DEVA:[JONES]SHCINI.INI when it starts up. If the initialization file is found, SHOW CLUSTER executes the procedure before beginning the display.

If you do not define SHOW_CLUSTER\$INIT or it does not include a directory specification, SHOW CLUSTER searches the current default directory for a file named SHOW_CLUSTER.INI.

VMScLuster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

2. Customize the display using SHOW CLUSTER commands during a continuous SHOW CLUSTER session.
3. Preserve the command sequence by entering the following command:

Command> SAVE SHOW_CLUSTER\$INIT.INI

You must specify SHOW_CLUSTER\$INIT.INI, because the SAVE command creates a file with a file type of .COM by default. SHOW CLUSTER looks for an .INI file when it searches for a startup initialization file.

You can edit the file that the SAVE command creates to include comments or to improve its efficiency. For more information, see the SAVE command in the *OpenVMS System Management Utilities Reference Manual*.

Instead of having SHOW CLUSTER build an initialization file, you can build one yourself in the same way you build a command procedure. The next section provides guidelines for creating a command procedure.

20.3.7 Using Command Procedures Containing SHOW CLUSTER Commands

You can create command procedures that contain SHOW CLUSTER commands. Such files let you modify display characteristics without having to enter commands interactively. You can use command procedures during a continuous SHOW CLUSTER session to perform a series of commands, for example, to customize the output of the display.

Following are guidelines for writing command procedures that contain SHOW CLUSTER commands:

- Use any valid SHOW CLUSTER commands.
- Nest command procedures up to 16 levels deep.
- Include the SHOW CLUSTER command INITIALIZE as the first command in the file. The INITIALIZE command ensures that the report is in a known state before any commands are executed to modify it.

Notes

Do not include an EXIT command at the end of the command procedure. The EXIT command terminates SHOW CLUSTER and erases the SHOW CLUSTER display before you can see it.

Also, do not run SHOW CLUSTER command procedures from a batch job.

The following command procedure customizes a report display:

```
!  
! Include only the node field from the default display; show votes  
! and quorum for each node and for the cluster as a whole.  
!  
INITIALIZE  
REMOVE SOFTWARE, STATUS  
ADD VOTES, QUORUM, CL_VOTES, CL_QUORUM
```

This command procedure removes the SOFTWARE and STATUS fields from the report and adds fields that provide information about the cluster quorum and votes.

VMScluster Considerations

20.3 Using the Show Cluster Utility (SHOW CLUSTER)

To execute a command procedure during a continuous SHOW CLUSTER session, specify the Execute Procedure (@) command, along with the file name of the command procedure. The default file type for command procedure files is .COM.

Example

The following command executes a command procedure named SYSMOD.COM:

```
Command> @SYSMOD
```

In this example, the default file type .COM is assumed because the file type is omitted.

For more information on creating command procedures, see the SAVE command in the *OpenVMS System Management Utilities Reference Manual*.

20.4 Understanding SYSMAN and VMScluster Management

The System Management utility (SYSMAN) provides two kinds of support for VMScluster management:

- Cluster-specific commands, CONFIGURATION SET and CONFIGURATION SHOW, that you can use to manage security data and system time in a VMScluster
- Access to DCL-level commands with the DO command, which gives you the ability to apply a single DCL command across an entire VMScluster, rather than having to enter the command on each node

Each SYSMAN command requires a specific level of privilege. For more information on each command, see the *OpenVMS System Management Utilities Reference Manual*.

20.5 Using SYSMAN to Manage Security and System Time

You can manage security data and system time for a VMScluster with the SYSMAN CONFIGURATION commands. Table 20-5 summarizes these CONFIGURATION commands and their functions.

Table 20-5 SYSMAN CONFIGURATION Commands

Command	Function
CONFIGURATION SET CLUSTER_AUTHORIZATION	Modifies the group number and password in a local area VMScluster
CONFIGURATION SHOW CLUSTER_AUTHORIZATION	Displays the group number and multicast address of a local area VMScluster
CONFIGURATION SET TIME	Updates system time
CONFIGURATION SHOW TIME	Displays current system time

20.5.1 Modifying the Group Number and Password

The group number identifies the group of nodes in the VMScluster, and the associated Ethernet address is used to send messages to all nodes in the cluster. The VMScluster password protects the integrity of the VMScluster membership.

VMScLuster Considerations

20.5 Using SYSMAN to Manage Security and System Time

Using the `CONFIGURATION SET CLUSTER_AUTHORIZATION` command modifies the group number and password, as recorded in `SYS$SYSTEM:CLUSTER_AUTHORIZE.DAT`. Normally, you do not need to alter records in the `CLUSTER_AUTHORIZE.DAT` file.

If your configuration has multiple system disks, SYSMAN automatically updates each copy of `CLUSTER_AUTHORIZE.DAT`, provided that you have defined the environment as a VMScLuster with the `SET ENVIRONMENT/CLUSTER` command.

Caution

If you change either the group number or password, you must reboot the entire VMScLuster.

You cannot display the VMScLuster password for security reasons, but you can display the group number and group multicast address with the `CONFIGURATION SHOW CLUSTER_AUTHORIZATION` command.

Examples

1. The following command example sets the environment to a specific cluster, sets privilege to `SYSPRV`, and modifies the VMScLuster password:

```
SYSMAN> SET ENVIRONMENT/CLUSTER/NODE=NODE21
SYSMAN> SET PROFILE/PRIVILEGE=SYSPRV
SYSMAN> CONFIGURATION SET CLUSTER_AUTHORIZATION/PASSWORD=GILLIAN
%SYSMAN-I-CAFOLDGROUP, existing group will not be changed
%SYSMAN-I-GRPNOCHG, Group number not changed
SYSMAN-I-CAFREBOOT, cluster authorization file updated.
The entire cluster should be rebooted.
```

2. The following command example displays the group number and multicast address for `NODE21`. Because the group number and password on other nodes in the VMScLuster are identical, no further information is displayed.

```
SYSMAN> CONFIGURATION SHOW CLUSTER_AUTHORIZATION
Node NODE21: Cluster group number 65240
Multicast address: AB-00-04-01-F2-FF
```

20.5.2 Modifying the System Time

Use the `CONFIGURATION SET TIME` command to modify system time for nodes in a VMScLuster, as well as for individual nodes. You can specify time values in the following format:

`[dd-mmm-yyyy[:]] [hh:mm:ss.cc]`

You can also enter delta time values. See the *OpenVMS User's Manual* for more information about time formats.

In a VMScLuster environment, SYSMAN sets the time on each node to the value you specify. However, if you do not specify a value, SYSMAN reads the clock on the node from which you are executing SYSMAN and assigns this value to all nodes in the VMScLuster. In a remote VMScLuster, SYSMAN reads the clock on the target node in the cluster and assigns that value to all nodes. Note that the time-of-year clock is optional for some processors; see your processor's hardware handbook for more information.

VMScLuster Considerations

20.5 Using SYSMAN to Manage Security and System Time

SYSMAN tries to ensure that all processors in the VMScLuster are set to the same time. Because of communication and processing delays, it is not possible to synchronize clocks exactly. However, the variation is typically less than a few hundredths of a second. If SYSMAN cannot set the time to within one-half second of the specified time, you receive a warning message that names the node that failed to respond quickly enough.

As a result of slight inaccuracies in each processor clock, times on various members of a VMScLuster tend to drift apart. The first two examples show how to synchronize system time in a VMScLuster.

Examples

1. The following procedure sets the time on all VMScLuster nodes to the value obtained from the local time-of-year clock, waits 6 hours, then resets the time for the VMScLuster:

```
$ SYNCH_CLOCKS:
$ RUN SYS$SYSTEM:SYSMAN
  SET ENVIRONMENT/CLUSTER
  CONFIGURATION SET TIME
  EXIT
$ WAIT 6:00:00
$ GOTO SYNCH_CLOCKS
```

2. The next example sets the environment to NODE21, NODE22, and NODE23, sets privilege, and modifies the system time on all three nodes:

```
SYSMAN> SET ENVIRONMENT/NODE=(NODE21,NODE22,NODE23)
SYSMAN> SET PROFILE/PRIVILEGE=LOG_IO
SYSMAN> CONFIGURATION SET TIME 12:38:00
```

3. The following example sets the environment to cluster and displays the system time for all nodes:

```
SYSMAN> SET ENVIRONMENT/CLUSTER/NODE=NODE23
SYSMAN> CONFIGURATION SHOW TIME
System time on node NODE21: 19-APR-1995 13:32:19.45
System time on node NODE22: 19-APR-1995 13:32:27.79
System time on node NODE23: 19-APR-1995 13:32:58.66
```

20.5.2.1 Resetting System Time After January 1

VAX

The Time of Day Register (TODR), which the system uses to maintain system time, has a limit of approximately 15 months. Between January 1 and April 1, reset the system time; otherwise, the following problems might occur:

- The first time in a new year that you reboot a VAXcluster system or a node in the system, one or more nodes display any of the following system times:
 - A year in the past
 - A year in the future, which might cause passwords to expire and other difficulties
 - A correct time, but a SHOW SYSTEM command indicates that the system has been up since a time in the 1800s
- Even if you correct the system time during system boot, the following problems might remain:
 - A SHOW SYSTEM command displays an incorrect up time such as a date in the 1800s
 - The error log report (ERRLOG) shows errors for a year in the future

VMScIuster Considerations

20.5 Using SYSMAN to Manage Security and System Time

- Batch jobs are waiting for a year in the future
- Files have a creation or modification date in the future

Because the TODR has an approximate limit of 15 months, the system maintains time by combining the TODR value with a base time recorded in the base system image (SYS\$LOADABLE_IMAGES:SYS.EXE). The definition of base time is:

01-JAN-CURRENT_YEAR 00:00:00.00

Because all TODRs ordinarily have the same base, multiple CPUs can boot off the same system disk, and you can use multiple system disks on one CPU; the system sets the time correctly.

When a SET TIME command is issued (with or without specifying a time), OpenVMS does the following:

1. Writes the current time to the system image file
2. Resets the TODR as an offset within the current year

In a VAXcluster system (or for a node that is not part of the cluster), when you set the time, the TODR and the base time in the system image are reset with the values for the new year. However, multiple systems might share the system image. This does not normally cause a problem except after the first day of a new year.

Note

The system issues the SET TIME command when it boots and as a part of the normal SHUTDOWN command procedure.

By December, each node has a very large offset stored in the TODR (from the base time of 1-JAN of that year). When the time advances to a new year, the system image still has the old year and the TODR values are still large.

After January 1, if a SET TIME command is issued on any node (or any node is shut down using SHUTDOWN.COM), the following happens:

1. The new year becomes the base year
2. The system resets the TODR on that node
3. The other nodes still have a large value in the TODR

After these three events occur, if a node that has a large TODR crashes and rejoins the cluster, its system time is initially in the next year (applying the large TODR to the new year). This system time is recorded as the system's boot time. When the node joins the cluster, its time is set to the correct value but the boot time remains one year in the future. Certain forms of the SHOW SYSTEM command compare current time to boot time; in this instance, SHOW SYSTEM displays incorrect values.

If a system disk is used at different times by different, unclustered CPUs or if different system disks are used at different times on the same CPU, the system might incorrectly set the time to a year in the future or a year in the past, depending on how the CPU's TODR and the value recorded on the system disk become unsynchronized:

- Sharing a system disk across multiple CPUs pushes the time into the future

VMScluster Considerations

20.5 Using SYSMAN to Manage Security and System Time

- Using multiple disks on one CPU pushes the time into the past

Example

The following example uses SYSMAN commands to reset the time on all nodes in a VAXcluster system:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> SET PROFILE/PRIVILEGE=(LOG_IO,SYSLCK)
SYSMAN> CONFIGURATION SET TIME 05-JUN-1994:12:00:00 1
SYSMAN> EXIT
```

Note

In a node that is not part of a VAXcluster system, use the SET TIME command and specify a time. If you do not specify a time, the SET TIME command updates the system time using the time in the TODR.

Note

If you are running the Digital Distributed Time Service (DECdts) on your system, you must use it to set the time.

20.6 Using the SYSMAN DO Command to Manage a VMScluster

Using the SYSMAN command DO enables you to execute a DCL command or command procedure on all nodes in the current environment. This is convenient when you are performing routine system management tasks on nodes in the VMScluster system, such as:

- Installing images
- Starting up software
- Checking devices
- Checking memory

Each DO command executes as an independent process, so there is no process context retained between DO commands. For this reason, you must express all DCL commands in a single command string, and you cannot run a program that expects input.

In a VMScluster environment, SYSMAN executes the commands sequentially on all nodes in the VMScluster. Each command executes completely before SYSMAN sends it to the next node in the environment. Any node that is unable to execute the command returns an error message. SYSMAN displays an error message if the timeout period expires before the node responds.

In a dual-architecture heterogeneous VMScluster running both OpenVMS VAX and OpenVMS Alpha, some uses of the DO command may require special handling. For example, if you are installing images that are named differently in each architecture, you can still use the DO command if you create logical name tables for VAX and for Alpha nodes. See the example sequence that follows this description for an example.

20.6 Using the SYSMAN DO Command to Manage a VMScluster

VMSccluster Considerations

Some DCL commands, such as MOUNT/CLUSTER or SET QUORUM/CLUSTER, operate clusterwide by design. It is best to avoid using these kinds of commands with the DO command in SYSMAN when the environment is set to cluster. As alternatives, you could leave SYSMAN temporarily with the SPAWN command and execute these commands in DCL, or you could define the environment to be a single node within the VMScluster.

Examples

1. The following example installs an image on a VMScluster. First, it adds CMKRNL and SYSPRV privileges to the current privileges because they are required by INSTALL and AUTHORIZE. The DO INSTALL command installs the file STATSHR. The DO MCR AUTHORIZE command sets up an account for user Jones, specifying a password and a default device and directory.

```
SYSMAN> SET PROFILE/PRIVILEGES=(CMKRNL,SYSPRV)/DEFAULT=SYS$SYSTEM
SYSMAN> DO INSTALL ADD/OPEN/SHARED WRKD$:[MAIN]STATSHR
SYSMAN> DO MCR AUTHORIZE ADD JONES/PASSWORD=COLUMBINE -
_SYSMAN> /DEVICE=WORK1/DIRECTORY=[JONES]
```

2. The following example sets the environment to cluster and starts up a software product called XYZ on each node in the VMScluster:

```
SYSMAN>SET ENVIRONMENT/CLUSTER
%SYSMAN-I-ENV, Current command environment:
      Clusterwide on local cluster
      Username SMITH will be used on nonlocal nodes
SYSMAN> DO @SYS$STARTUP:XYZ_STARTUP
```

3. The following example shows how you can define logical names for VAX and Alpha nodes in a dual-architecture heterogeneous VMScluster, so that you can use the DO command to install architecture-specific images.

```
$ CREATE/NAME_TABLE/PARENT=LN$SYSTEM_DIRECTORY SYSMAN$NODE_TABLE
$ DEFINE/TABLE=SYSMAN$NODE_TABLE ALPHA_NODES NODE21,NODE22,NODE23
$ DEFINE/TABLE=SYSMAN$NODE_TABLE VAX_NODES NODE24,NODE25,NODE26
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/NODE=ALPHA_NODES
%SYSMAN-I-ENV, current command environment:
      Individual nodes: NODE21,NODE22,NODE23
      Username BOUCHARD will be used on nonlocal nodes

SYSMAN> DO INSTALL REPLACE SYS$LIBRARY:DCLTABLES.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE21
%SYSMAN-I-OUTPUT, command execution on node NODE22
%SYSMAN-I-OUTPUT, command execution on node NODE23
SYSMAN> DO INSTALL REPLACE SYS$SYSTEM:DEC_FORTRAN.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE21
%SYSMAN-I-OUTPUT, command execution on node NODE22
%SYSMAN-I-OUTPUT, command execution on node NODE23

SYSMAN> SET ENVIRONMENT/NODE=VAX_NODES
%SYSMAN-I-ENV, current command environment:
      Individual nodes: NODE24,NODE25,NODE26
      Username BOUCHARD will be used on nonlocal nodes

SYSMAN> DO INSTALL REPLACE SYS$LIBRARY:DCLTABLES.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE24
%SYSMAN-I-OUTPUT, command execution on node NODE25
%SYSMAN-I-OUTPUT, command execution on node NODE26
SYSMAN> DO INSTALL REPLACE SYS$SYSTEM:FORTAN$MAIN.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE24
%SYSMAN-I-OUTPUT, command execution on node NODE25
%SYSMAN-I-OUTPUT, command execution on node NODE26
```


VMSccluster Considerations

20.6 Using the SYSMAN DO Command to Manage a VMSccluster

4. The following example shows which files are open on DISK2. You might use this if you want to dismount DISK2 and need to see which users in the VMSccluster have files open.

```
SYSMAN >SET ENVIRONMENT/CLUSTER
%SYSMAN-I-ENV, Current command environment:
    Clusterwide on local cluster
    Username SMITH will be used on nonlocal nodes
SYSMAN> DO SHOW DEVICE/FILES DISK2:

%SYSMAN-I-OUTPUT, command execution on node NODE21
Files accessed on device $1$DIA2: (DISK2, NODE22) on 14-MAY-1995 15:44:06.05
Process name      PID      File name
                  00000000  [000000]INDEXF.SYS;1

%SYSMAN-I-OUTPUT, command execution on node NODE22
Files accessed on device $1$DIA2: (DISK2, NODE21) on 14-MAY-1995 15:44:26.93
Process name      PID      File name
                  00000000  [000000]INDEXF.SYS;1

%SYSMAN-I-OUTPUT, command execution on node NODE23
Files accessed on device $1$DIA2: (NODE21, NODE22) on 14-MAY-1995 15:45:01.43
Process name      PID      File name
                  00000000  [000000]INDEXF.SYS;1

%SYSMAN-I-OUTPUT, command execution on node NODE24
Files accessed on device $1$DIA2: (NODE22, NODE21) on 14-MAY-1995 15:44:31.30
Process name      PID      File name
                  00000000  [000000]INDEXF.SYS;1

Susan Scott      21400059  [SCOTT]DECW$SM.LOG;228
_FTA7:           214000DD  [SCOTT]CARE_SDML.TPU$JOURNAL;1

%SYSMAN-I-OUTPUT, command execution on node NODE25
Files accessed on device $1$DIA2: (NODE21, NODE22) on 14-MAY-1995 15:44:35.50
Process name      PID      File name
                  00000000  [000000]INDEXF.SYS;1
DECW$SESSION      226000E6  [SNOW]DECW$SM.LOG;6
_FTA17:           2260009C  [SNOW.MAIL]MAIL.MAI;1
SNOW_1            2260012F  [SNOW.MAIL]MAIL.MAI;1
SNOW_2            22600142  [SNOW.MAIL]MAIL.MAI;1
SNOW_3            22600143  [SNOW.MAIL]MAIL.MAI;1
```


VMScluster Considerations

20.6 Using the SYSMAN DO Command to Manage a VMScluster

5. The following example shows how much memory is available on the nodes in a VMScluster. You might use this if you are installing software and want to know if each node has enough memory available.

```

SYSMAN > SET ENVIRONMENT/NODE=(NODE21,NODE22)
%SYSMAN-I-ENV, Current command environment:
Clusterwide on local cluster
Username SMITH will be used on nonlocal nodes
SYSMAN> DO SHOW MEMORY
%SYSMAN-I-OUTPUT, command execution on node NODE21
System Memory Resources on 14-MAY-1995 15:59:21.61
Physical Memory Usage (pages):      Total      Free      In Use      Modified
Main Memory (64.00Mb)              131072     63955     65201      1916
Slot Usage (slots):                 Total      Free      Resident    Swapped
Process Entry Slots                  360        296         64          0
Balance Set Slots                    324        262         62          0
Fixed-Size Pool Areas (packets):     Total      Free      In Use      Size
Small Packet (SRP) List              10568     1703     8865       128
I/O Request Packet (IRP) List        3752      925     2827       176
Large Packet (LRP) List              157        28      129      1856
Dynamic Memory Usage (bytes):         Total      Free      In Use      Largest
Nonpaged Dynamic Memory              1300480    97120    1203360    60112
Paged Dynamic Memory                 1524736   510496    1014240   505408
Paging File Usage (pages):           Free      Reservable    Total
DISK$MTWAIN_SYS:[SYS0.SYSEXE]SWAPFILE.SYS
                                     10000      10000      10000
DISK$MTWAIN_SYS:[SYS0.SYSEXE]PAGEFILE.SYS
                                     60502     -52278     100000
Of the physical pages in use, 19018 pages are permanently allocated to VMS.

%SYSMAN-I-OUTPUT, command execution on node NODE22
System Memory Resources on 14-MAY-1995 15:59:42.65
Physical Memory Usage (pages):      Total      Free      In Use      Modified
Main Memory (32.00Mb)              65536     44409     20461      666
Slot Usage (slots):                 Total      Free      Resident    Swapped
Process Entry Slots                  240        216         24          0
Balance Set Slots                    212        190         22          0
Fixed-Size Pool Areas (packets):     Total      Free      In Use      Size
Small Packet (SRP) List              5080     2610     2470       128
I/O Request Packet (IRP) List        3101     1263     1838       176
Large Packet (LRP) List              87         60         27      1856
Dynamic Memory Usage (bytes):         Total      Free      In Use      Largest
Nonpaged Dynamic Memory              1165312   156256    1009056   114432
Paged Dynamic Memory                 1068032   357424     710608   352368
Paging File Usage (pages):           Free      Reservable    Total
DISK$MTWAIN_SYS:[SYS1.SYSEXE]SWAPFILE.SYS
                                     10000      10000      10000
DISK$MTWAIN_SYS:[SYS1.SYSEXE]PAGEFILE.SYS
                                     110591     68443     120000
Of the physical pages in use, 9056 pages are permanently allocated to VMS.

```


Network Considerations

On OpenVMS systems, two types of DECnet functionality are available:

- DECnet for OpenVMS, which is the default version of DECnet that ships with the OpenVMS operating system.
- DECnet/OSI, which is a separately installable family of products that enable OpenVMS operating systems to communicate with each other and with systems produced by other vendors. DECnet/OSI includes Digital's implementation of the following:
 - The Open Systems Interconnection (OSI) communications specifications, as defined by the International Organization for Standardization (ISO).
 - Digital's communications architecture, Digital network architecture (DNA) Phase V, which is also backward compatible with the Phase IV architecture.

For an introduction to DECnet/OSI, see *DECnet/OSI for OpenVMS Introduction and Planning*.

Note

With the exception of full names, which is a DECnet/OSI feature, this chapter describes only DECnet for OpenVMS functionality.

You can connect your system to a network by means of the DECnet interface. With this interface, you can link computers into flexible configurations to exchange information, share resources, and perform distributed processing.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Assigning node names	Section 21.1
Providing security for your node	Section 21.4
Accessing remote services	Section 21.5
Providing host services	Section 21.6.1
Monitoring the network	Section 21.6.2
Testing the network	Section 21.6.3
Shutting down and restarting the network	Section 21.6.4

Network Considerations

This chapter explains the following concepts:

Concept	Section
A DECnet for OpenVMS network	Section 21.2
How an OpenVMS system can be part of a network	Section 21.2.1
How nodes are connected to the network	Section 21.2.2
How a SCSI bus connects multiple nodes	Section 21.2.3
The configuration database	Section 21.2.4
How your system becomes a node in the network	Section 21.2.5
Preparations for joining a network	Section 21.3

For more details, refer to the following manuals:

Manual	Description
<i>DECnet for OpenVMS Guide to Networking</i>	Provides an introduction to networking on the system.
<i>DECnet for OpenVMS Networking Manual</i>	Includes conceptual and usage information.
<i>DECnet for OpenVMS Network Management Utilities</i>	Explains how to use the Network Control Program (NCP) utility.

Where appropriate, this chapter refers to specific manuals in this group.

21.1 Assigning Node Names

Naming conventions for DECnet node names correspond to the two types of DECnet functionality:

- DECnet for OpenVMS node names
These names are used in the default version of DECnet shipped with the OpenVMS operating system. Refer to the *DECnet for OpenVMS Networking Manual* for more information.
- DECnet/OSI full names
Full names are hierarchically structured DECnet node names that can be stored in a DECdns naming service. Full names can be a maximum of 255 bytes long.

21.1.1 Syntax for Full Names

Full names have the following general syntax:

namespace:.directorydirectory.node-name

where:

<i>namespace</i>	Identifies the global naming service
<i>directorydirectory</i>	Defines the hierarchical directory path within the naming service
<i>node-name</i>	Is the specific object defining the DECnet node

The node full name must begin with the namespace, followed by a colon (:). The directory path must begin with a period (.).

Example

In the following example of a full name, OMNI is the name of a namespace in a global naming service and .US.Massachusetts.Boston is the directory path to the object, RUBY, which represents the node:

OMNI:.US.Massachusetts.Boston.RUBY

The system stores a full name as you enter it, preserving uppercase and lowercase entries. However, when matching an entry with a stored full name, the system is case insensitive; in other words, if the user enters Omni, the system recognizes it as equivalent to OMNI.

21.1.2 Considerations for Assigning Full Names

Consider the following when assigning node names:

- Node full names must be unique within your network.
- OpenVMS VAX systems support any type of node name except one containing an odd number of double quotes.
- You must place quotes (" ") around a node name if the node name:
 - Contains a space, tab, comma, left parenthesis, right parenthesis, single quote ('), double quote ("), slash (/), exclamation point (!), plus sign (+), or at sign (@)
 - Contains the character sequence of two colons (::)
 - Starts or ends with a colon (:)
- If a DECnet node name contains double quotes, you must duplicate each double quote. Also, you must be sure to pair double quotes within a full name. For example, the node name *foo:"bar"* must appear as *foo:""bar""*.

Although OpenVMS enforces few rules on the syntax of node names, the DECnet for OpenVMS software running on your system limits the actual set of valid node names you can use. ♦

For more information on full names, refer to your DECnet/OSI documentation.

21.1.3 Setting Up a Node Name Strategy

As network manager, you are allowed to assign some or all of the names related to clusters and networking. You can, if you like, assign a different name to each object in the system. However, such an approach might easily confuse your users.

Digital recommends that you establish a system for assigning names. In the long run, assigning names in a systematic way will save you and your users time and trouble. Table 21-1 is an example of the type of methodology you might want to use:

Table 21-1 Sample System Names

Name for...	Example
DECnet/OSI full name	OMNI:.US.Massachusetts.Boston.RUBY
DECnet/OSI node synonym	RUBY

(continued on next page)

Network Considerations

21.1 Assigning Node Names

Table 21-1 (Cont.) Sample System Names

Name for...	Example
DECnet for OpenVMS node name	RUBY
SCS node name	RUBY
LAT service name	RUBY
TCP/IP name	ruby.omni.com

21.2 Understanding DECnet for OpenVMS Networks

A **network** is a means of connecting computers that allows them to share or transfer information or communications. A network includes two or more computers that are connected, and the hardware and software that make those connections.

DECnet for OpenVMS is the name of the software and hardware products that, collectively, provide the means for various Digital operating systems to participate in a network. DECnet allows an OpenVMS operating system to function as a network node. As a part of a network, an OpenVMS system can communicate with all types of OpenVMS systems, as well as with many systems that are not OpenVMS and that support DECnet.

Table 21-2 defines terms related to DECnet networks.

Table 21-2 DECnet for OpenVMS Network Terminology

Term	Definition
Node	A computer system that is connected to another system in a network—by means of cables, telephone lines, microwave and satellite links, for example. An adjacent node is one that is connected to your node by a single physical line.
Line	Physical data path that connects adjacent nodes in a network. A communications line connects your computer to the DECnet network.
Circuit	Communications data path that connects adjacent nodes in a network. A circuit is not a physical data path but, rather, a logical connection that operates over a physical connection (a line). All input and output (I/O) between nodes takes place over circuits. You can configure a node to have a number of active circuits and lines that connect it to other systems in the network.
Logical link	Connects two processes and carries a stream of two-way communications traffic between the processes over a circuit. A single circuit established between two nodes can support many logical links concurrently.
Object	Process to which the logical link connects. Some objects are DECnet system programs—for example, the MAIL object; other objects are user-written programs. For two programs to communicate over the network, the source program on the local node establishes a logical link with the object on the remote node.

(continued on next page)

Network Considerations


21.2 Understanding DECnet for OpenVMS Networks

Table 21-2 (Cont.) DECnet for OpenVMS Network Terminology

Term	Definition
Ethernet	A single shared network channel, with all nodes having equal access to the channel. Ethernet offers local and remote connections as one integral network.

Figure 21-1 shows lines and circuits connecting nodes in a DECnet network.

Figure 21-1 Network Nodes, Circuits, and Lines



A DECnet network is decentralized. Many nodes connected to the network can communicate with each other without having to go through a central node. As a member of a multinode network, a node can communicate with any other network node, not merely with those nodes physically attached to it. This feature allows users to gain access to software facilities that might not exist on their particular nodes.

DECnet Routing

In a network of more than two nodes, the process of directing a data message from a source node to a destination node is called **routing**. DECnet supports **adaptive routing**, which routes messages through the network over the most cost-effective path. Adaptive routing also reroutes messages automatically if a circuit becomes disabled or a lower-cost path becomes available.

Nodes can be either routing nodes (called **routers**) or nonrouting nodes (known as **end nodes**). Both routers and end nodes can send messages to and receive messages from other nodes in the network. Following are the differences between a router and an end node.

- **Router**

Routing node; has the ability to forward or route messages from itself to another node.

A routing node can serve as an intermediate node on a path between two nodes exchanging messages, if the two nodes have no direct physical link to each other. Any node that has two or more active circuits connecting it to the network must be a router.

DECnet supports routing within each area; DECnet also supports a second, higher level of routing that links the areas, resulting in less routing traffic throughout the network.

The higher levels of routing are the following:

- **Level 1 routers**

These are nodes that perform routing within a single area.

Network Considerations

21.2 Understanding DECnet for OpenVMS Networks

Alpha

- **Level 2 routers (or area routers)**

These are nodes that perform routing between areas as well as within their own area.

On Alpha systems, DECnet does not support routing. This end-node only (nonrouting) capability means that an Alpha node can receive packets addressed to it and can send packets to other nodes, but it cannot route packets. For more information on DECnet restrictions on Alpha systems, refer to *A Comparison of System Management on OpenVMS AXP and OpenVMS VAX*. ♦

- **End node**


Nonrouting node; can have only one active circuit connecting it to the network.

DECnet Configurations

DECnet supports configurations for both large and small networks. A typical small network might consist of two to four nodes. A maximum of 1023 nodes is possible in an undivided network, but the optimum number is approximately 200 to 300 nodes, depending on the topology.

Figure 21-2 illustrates a small Ethernet configuration of four nodes. Three VAXstation-based end nodes and one router (the VAX-9000) are connected to the Ethernet.

Figure 21-2 Example of a Small Local Area Network Configuration



A DECnet network has built-in flexibility in topology and performance. Its architecture adheres to industry standards, and is designed to permit easy expansion and incorporation of new developments in data communications. DECnet offers the option of communicating over different kinds of network connections, which are, for the most part, transparent to the general user of the network.

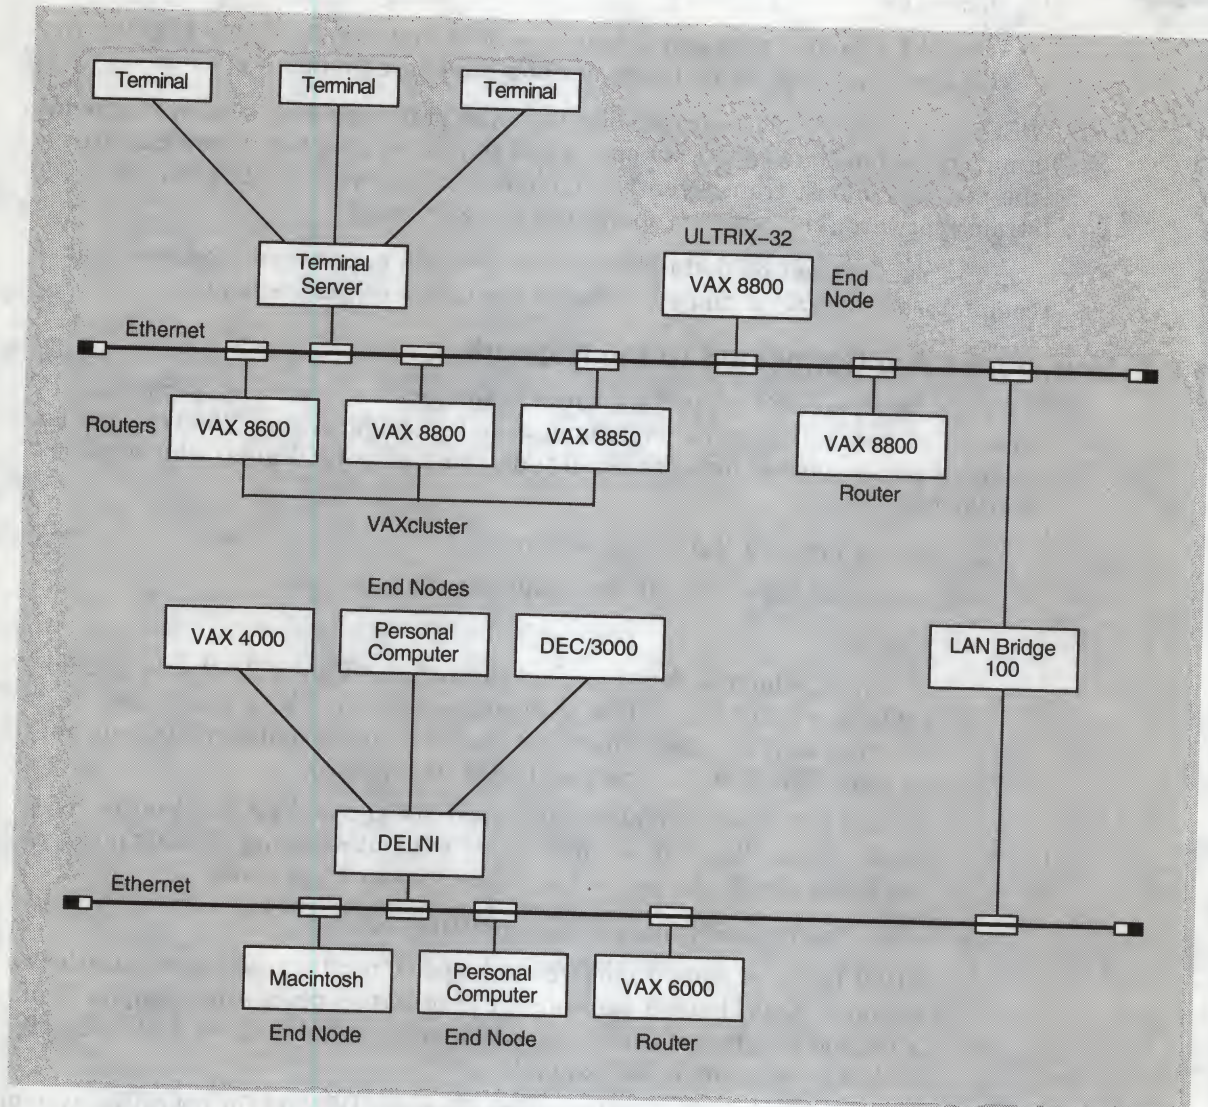
You can divide very large DECnet networks into multiple areas: up to 63 areas, each containing a maximum of 1023 nodes. In a multiple-area network, nodes are grouped into separate areas, with each area functioning as a subnetwork. Nodes in any area can communicate with nodes in other areas.

Figure 21-3 is an example of a large local area DECnet configuration that illustrates a variety of ways in which you can connect OpenVMS systems to the network. The figure indicates whether a particular node is a router or an end node.

Network Considerations

21.2 Understanding DECnet for OpenVMS Networks

Figure 21-3 Large Local Area Network Configuration



ZK-6358-GE

Figure 21-3 shows a larger local area network (LAN) configuration in which two Ethernets are connected by a LAN bridge. Various kinds of operating systems, including nodes in a cluster, are connected directly to the Ethernet. In the figure, a group of small systems is connected to the Ethernet by means of a DELNI. Individual terminal users can gain access to Ethernet nodes through a terminal server.

21.2.1 How an OpenVMS System Can Be Part of a Network

As the OpenVMS network interface, DECnet supports both the protocols necessary for communicating over the network and the functions necessary for configuring, controlling, and monitoring the network.

You can configure DECnet networking software on any OpenVMS operating system. A DECnet node can communicate with the following:

- Other DECnet nodes in the network
- Nodes with any other operating system that supports DECnet

VAX

- On VAX systems, nodes on other networks, by means of packet-switching networks
- On VAX systems, nodes with foreign vendor systems, by means of gateways, bridges, and other special software and hardware products ♦

DECnet is completely integrated into the OpenVMS operating system; it provides a natural extension of local I/O operations to remote systems. Users can use the network almost transparently. Implementing network applications is straightforward, and network operations are efficient.

You can use DECnet on a standalone node—to run application programs that communicate directly with each other at the task level, for example.

21.2.2 How Nodes Are Connected to the Network

DECnet for OpenVMS supports a variety of network connections, permitting you to link computers and terminals in flexible configurations. The type you use depends on the type of network connection you make: local area, wide area, or worldwide:

- Local area network (LAN) connections

For local area networks, DECnet supports the following:

– Ethernet

Ethernet, which is shown in Figure 21-1 and Figure 21-2, is a coaxial cable to which each system or device is connected by a single line. In an office or other area where personal computers and workstations are located, ThinWire Ethernet cabling is usually used.

On the Ethernet, a single, shared network channel LAN, all nodes have equal access. You can add new nodes without affecting existing nodes on the Ethernet. An Ethernet can support up to 1,023 nodes.

– Fiber Distributed Data Interface (FDDI) LANs

FDDI LANs provide a reliable, high-speed, multiaccess communications channel. This channel can connect information processing equipment in a limited geographic area, such as an office, a building, or a complex of buildings—a campus, for example.

VAX

On VAX systems, nodes in a VAXcluster require DECnet for operating system connections. Each node in a cluster can be connected to an Ethernet that provides the data link for the cluster. If an Ethernet is not available, you can configure the CI computer interconnect used by the VAXcluster to be the data link between the cluster nodes. FDDI LANs also support VAXcluster technology and let you configure a computer system with its components spread out over several miles. ♦

- Wide area network (WAN) connections

VAX

On VAX systems, DECnet offers comprehensive wide area network support and long-haul connectivity over point-to-point and multipoint connections:

– Point-to-point connections

These use DDCMP, Digital's data communications message protocol, and are synchronous or asynchronous:

- * Synchronous devices provide high-speed connections over local lines or telephone lines (using modems).

Network Considerations

21.2 Understanding DECnet for OpenVMS Networks

- * Asynchronous devices provide low-speed, low-cost connections over terminal lines that are switched on for network use either permanently (a static connection) or temporarily (a dynamic connection). For example, a user at a MicroVAX terminal can configure a dialup line to another computer as a dynamic asynchronous DECnet line for the duration of a call.
- Multipoint connections

These consist of two or more nodes connected by a synchronous DDCMP communications channel, with one node controlling the channel.
- Multiple-site VMScluster systems

VMScluster systems support DS3 technology, also called T3. Using DS3 as an interconnect, nodes in a VMScluster system can be located in multiple, geographically separate sites as far apart as 150 cable miles (the physical cable distance, not the physical mile distance).
- Worldwide network connections

DECnet supports worldwide communications with a range of different networks through packet switching networks and gateways. ♦

21.2.3 Connecting Multiple Nodes to a SCSI Bus

Alpha


One of the benefits of VMScluster systems is that multiple computers can simultaneously access storage devices connected to a VMScluster storage interconnect. When multiple Alpha nodes in a VMScluster connect to a single Small Computer Systems Interface (SCSI) bus they can share access to SCSI storage devices directly. This capability allows you to build highly available servers for shared access to SCSI storage.

Figure 21-4 shows a VMScluster configuration that uses a SCSI interconnect for shared access to SCSI devices. Note that another interconnect (for example, a local area network [LAN]) is required for node-to-node VMScluster (System Communications Architecture [SCA]) communications. For more information, see the *OpenVMS Version 6.2 New Features Manual*.

Network Considerations

21.2 Understanding DECnet for OpenVMS Networks

Figure 21-4 Highly Available Servers for Shared SCSI Access



21.2.4 Understanding the Configuration Database

The system manager at each node in the network is responsible for the DECnet for OpenVMS **configuration database** for the node. Each node in the network has a configuration database, which is stored in the SYS\$SYSTEM:NETNODE_REMOTE.DAT file. You can change the location of this file by defining the logical name NETNODE_REMOTE in the SYLOGICALS.COM file. (See Section 5.2.5 for details.)

Besides storing information about other nodes in the network with which the local node can communicate, a configuration database contains the following information:

- Files that describe the following:
 - The local (executor) node
 - The circuits and lines that connect the local node to the network
- Information on the logging collection points (such as the logging monitor) to which network events are reported
- Object databases that describe objects (such as MAIL) known to the network

Network Considerations

21.2 Understanding DECnet for OpenVMS Networks

As system manager, you provide network component information, from the point of view of the local node, in the configuration database at the local node. You can use the Network Control Program (NCP) to build the network configuration database manually or to modify its contents. If you are configuring a node for the first time, you can use the automatic configuration command procedure, NETCONFIG.COM, to establish parameters needed to get DECnet running.

The configuration database is made up of a **volatile database** and a **permanent database**. Table 21-3 describes the two types of databases in the configuration database, compares the duration of changes you make to each, and specifies the NCP commands you use to specify database contents.

Table 21-3 Comparison of Volatile and Permanent Databases

Type of Database	Description	Effect of Modifications	NCP Commands You Can Use
Volatile database	Working copy of the database that reflects current network conditions	Changes exist only while the network is running	Use SET commands to specify the contents of the volatile database. Use CLEAR commands to delete or reset volatile database entries. OPER privilege is required to change a volatile database.
Permanent database	Provides the initial values for the volatile database when you start the network	Changes remain after the network is shut down, but do not affect the current system	Use DEFINE commands to establish the contents of the permanent database. Use PURGE commands to delete permanent database entries. SYSPRV privilege is required to change a permanent database.

21.2.5 How Your System Becomes a Node in the Network

As manager of a DECnet node, you are responsible for establishing your operating system as a node in the network. To do this, follow these steps. Subsequent sections explain these steps in more detail. For specific instructions for performing each step, see the *DECnet for OpenVMS Guide to Networking*.

1. Prepare your system, which includes:
 - a. Connecting the hardware
 - b. Planning how you want to configure your system
2. Make necessary purchases and registrations, including:
 - a. Purchasing a DECnet for OpenVMS license and Product Authorization Kit (PAK)
 - b. Using the License Management utility to register the PAK
3. Configure your node in the network, which includes:
 - a. Configuring your network environment automatically or manually
 - b. On VAX systems, optionally establishing asynchronous connections to other systems ♦
 - c. Verifying that your node is connected to the network
 - d. Providing security for your node

VAX

Network Considerations

21.3 Preparations for Joining a Network

21.3 Preparations for Joining a Network

This section describes preparations for connecting your system to an existing network. Specific instructions for performing these operations are in the *DECnet for OpenVMS Networking Manual*.

Operation	Description
Connect the hardware	To join the network and communicate with other systems, your system must have communications lines. (A communications line connects your computer to the DECnet network.)
Plan the configuration of your node	Planning the configuration of your node in an existing network usually involves coordinating with the system managers of other nodes in the network or with the manager of the network to ensure uniform assumptions about network parameter settings.
Purchase licenses and register a PAK	Before you can bring up your system as a node in the network, you must have a DECnet license and register a DECnet PAK on your system
Configure your node in the network	You can configure the node manually or automatically. You use the manual procedure if you want to modify an existing configuration. You use the automatic configuration procedure, <code>SY\$MANAGER:NETCONFIG.COM</code> , when you first join the network or when you reconfigure your node completely.
Verify your successful connection to the network	To verify your connection to the network, you can perform a number of tests that demonstrate whether your node can communicate with an adjacent node—that is, a node connected to your node by a single physical line. You can also use the DECnet Test Sender/DECnet Test Receiver (DTS/DTR) utility to test this connection.

21.4 Providing Security for Your Node

As manager of a network node, you can protect your system against unauthorized access by users on other nodes in the network by setting passwords for any accounts you create. You can also use the following security measures:

- Protect files and use proxy accounts

You use the DCL command `SET PROTECTION` to set limits on who can access the files in your account. If your file is protected, a user on a remote node must be able to specify the user name and password of a local account that has the appropriate privileges to access the file.

You can permit selected outside users to access particular accounts on your system without sending any explicit access control information over the network. You do this by creating a proxy account that allows a remote user to have access privileges on your node without having a private account on your node.

- Control access to your node

You can control access to the local node on two levels:

- Node level

To control the establishment of logical links with remote nodes, you can specify parameters in your network database access control; these parameters indicate which of the following logical links connections are permitted: `INCOMING`, `OUTGOING`, `BOTH`, or `NONE`.

Network Considerations

21.4 Providing Security for Your Node

To exclude unknown nodes, set Executor Default Access to NONE, thereby controlling the default.

- System level

When a remote user requests access to an object on the local node, a number of means of authorization are checked, including the following:

- * Is an explicit access control string available?
- * Does the user have a proxy account on the local node?
- * Does a default access account exist for the object at the local node?
- * Does a default nonprivileged DECnet account exist on the local node?

VAX

On VAX systems, you can also control access to the local node on non-broadcast circuits by using circuit-level access control. ♦

21.5 OpenVMS Support for TCP/IP Networking

OpenVMS supports Transmission Control Protocol/Internet Protocol (TCP/IP) parameters and qualifiers for the DCL commands SET HOST, COPY, and DIRECTORY. These commands invoke the TCP/IP layered software products that perform the following operations:

- Remote terminal service
- Remote file access
- Remote directory listings

For more detailed information, see the *TCP/IP Networking on OpenVMS Systems* manual. It introduces TCP/IP networking and the Internet, describes commonly used TCP/IP applications, and specifies formats of the DCL TCP/IP commands.

21.5.1 Remote Terminal Service

Remote terminal service operations enable a user at a local host (node) to interactively log in to a remote host. During this session, the local terminal operates as a virtual terminal on the remote host. OpenVMS clients running TCP/IP software can use the following SET HOST commands to access remote services:

DCL Command	Description
SET HOST/RLOGIN	Logs the user in to a remote host from a local host by using the Berkeley standard remote login virtual terminal protocol.
SET HOST/TELNET	Logs the user in to a remote host from a local host.
SET HOST/TN3270	Logs the user in to a remote IBM mainframe host from a local host. When TN3270 mode is active, the local keyboard emulates the keyboard of an IBM 3270 class terminal.

Both the RLOGIN and TELNET qualifiers perform remote login operations but implement them differently.

Network Considerations

21.5 OpenVMS Support for TCP/IP Networking

Example

The following command connects the local host to the remote host. The /AUTHENTICATE qualifier specifies that the remote host verify the user's identity so login can occur.

```
$ SET HOST/TELNET/AUTHENTICATE remotehst2
```

21.5.2 Remote File Access

Remote file access operations enable a user on a local host to copy files to and from a remote host, even though the file systems might be different as is the case between a host that uses OpenVMS and a host that does not. The following COPY commands download or upload files over the Internet.

DCL Command	Description
COPY/FTP	Copies files to or from a remote system by using the file transfer protocol.
COPY/RCP	Copies files to or from a remote system by using the Berkely standard remote copy protocol.

In most cases, transferring files to and from a remote system requires that the user have an account and password on that system. However, many computers on the Internet provide some type of public access, permitting users to log in to a special guest account. Some systems use the "anonymous FTP" service, which accepts a user name of *anonymous* and no password.

Example

The following command uses the /ANONYMOUS access qualifier to transfer a local ASCII text file to a remote system:

```
$ COPY/FTP/ASCII/ANON ovms_file1.c remotehst5::"/public/ovms_file2.c"
```

21.5.3 Remote Directory Listings

A user at a local host can list the directories of a remote system by using the DIRECTORY/FTP command. This is useful when transferring files to or from the remote system.

Example

The following command uses anonymous access to list the contents of the remote directory on the local host:

```
$ DIR/FTP/ANON remotehst6"Jones jpw"::"usr/public"
```

21.6 Managing a Network Node

Managing a network node usually requires regular monitoring to detect patterns of usage and error conditions on the network, and performing remote configuration of the network to control traffic patterns and accommodate network growth. You can perform maintenance procedures to prevent serious problems from developing, and troubleshooting procedures to resolve problems quickly.

The following sections briefly describe host services you might need to perform, software tools you can use to monitor and manage your DECnet network node, and instructions for shutting down and restarting the network. Refer to the *DECnet for OpenVMS Guide to Networking* for instructions for using these tools and the *DECnet for OpenVMS Networking Manual* for complete information on maintaining, controlling, testing, shutting down, and restarting the network.

21.6.1 Providing Host Services

As manager of a network node, you might also be called upon to provide DECnet host services for other nodes. Host services include:

- Loading system images and programs downline to unattended remote nodes
- Receiving for interpretation upline dumps of system images from nodes that have crashed

For example, DECnet permits you to load an operating system image or a terminal server image downline to a target node. Another DECnet host service involves connecting to an unattended remote node (for example, a diskless communications server) to act as its console.

21.6.2 Monitoring the Network

Using network tools, you can obtain statistics on network usage and routing parameters. Network logging files provide error statistics useful in diagnosing potential problems. Network Control Program (NCP) commands display the status of nodes, lines, and circuits in the network.

After collecting information about network activity, you can analyze the data you collect to determine whether the network is running properly and whether you should make changes to resolve problems or improve performance. Table 21-4 shows some of the ways you can monitor the network.

Table 21-4 Ways to Monitor the Network

Method	Use
NCP display commands	To determine the status and characteristics of components in the network
NCP counters	To obtain error and performance statistics on current network operations
Network events logged by DECnet	To report events to you as they happen
Other software tools, such as Ethernet configurator and the DECnet Test Sender/DECnet Test Receiver (DTS/DTR) utility	To learn more about network operation

21.6.2.1 Using NCP Display Commands

You can use the NCP commands SHOW and LIST to monitor network activity:

Command	Description
SHOW	These commands show the current condition of network components (no privileges required). Use these commands to monitor operation of the running network.
LIST	These commands list startup values assigned to network components (SYSPRV privilege required).

Table 21-5 shows some of the specific SHOW and LIST commands you can use and the information they display.

Network Considerations

21.6 Managing a Network Node

Table 21-5 NCP SHOW and LIST Commands

Command	Information Displayed
CHARACTERISTICS	Static information that does not normally change during network operations, such as the identification of the local node
COUNTERS	Counter information about circuits, lines, remote nodes, and the local node
EVENTS	Which network events are currently being logged to which logging collection point
LOGGING	Range of network events being logged by the DECnet Event Logging facility
STATUS	Dynamic information that usually indicates network operation for the running network, such as operational state of the local node
SUMMARY	Only the most useful information from both static and dynamic sources (the default)

21.6.2.2 NCP Counters

You can use NCP commands to display error and performance statistics about network components; you can do this at any time while the network is running. DECnet software uses counters to collect statistics automatically for the following:

- Executor node
- Remote nodes
- Circuits
- Lines

To display the contents of counters, you use NCP SHOW COUNTER commands. Following are typical examples of the commands:

```
$ RUN SYS$SYSTEM:NCP
NCP> SHOW EXECUTOR COUNTERS
NCP> SHOW NODE node-id COUNTERS
NCP> SHOW KNOWN CIRCUITS COUNTERS
NCP> SHOW KNOWN LINES COUNTERS
NCP> SHOW LINE line-id COUNTERS
NCP> EXIT
```

For the local node and remote nodes, counter statistics cover connection requests, user data traffic, timeouts, and errors. Specialized counters cover the following:

- Circuit counters: transmission of data packets over the circuit, timeouts, and errors.
- Line counters: transmission of bytes and data blocks over the line and relevant errors.

For a detailed explanation of NCP counters, see the *DECnet for OpenVMS Guide to Networking*. For a complete summary description of all network counters, including the probable causes of particular types of occurrences, refer to the *DECnet for OpenVMS Network Management Utilities*.

21.6.2.3 Using DECnet Event Logging

You can use the DECnet Event Logging facility to monitor important network events, including:

- Changes in circuit and line states (for example, a circuit failure)
- A node becoming reachable or unreachable
- Circuit and node counter values, logged before the counter is automatically set to 0
- Errors in data transmission
- User of invalid data link passwords

21.6.2.4 Using Other Software Tools

Table 21-6 shows some of the additional software tools that are available to view network activity or exercise network operations.

Table 21-6 Network Monitoring Tools

Tool	Description
NCP Ethernet configurator	Permits you to obtain a list of all systems on an Ethernet circuit or circuits
DECnet Test Sender/DECnet Test Receiver (DTS/DTR)	Cooperating tasks that perform various functions to exercise network task-to-task capabilities
Monitor utility	Monitors DECnet, displaying information about the use of system resources

21.6.3 Testing the Network

You can use the Network Control Program utility (NCP) to perform a series of tests to help determine whether the network is operating properly. These tests, which are called **loopback tests**, repeatedly send data through various network components that return the data to its source. If data is not looped successfully, or if the data is returned in a corrupted state, an NCP display indicates that the test failed; the display includes the reasons for the failure and the number of data messages not looped.

You can perform loopback tests at two levels:

Level	Description
Node	These loopback tests check the operation of logical links, routing, and other software.
Circuit	These loopback tests evaluate the operation of circuits. (You cannot perform these tests on asynchronous circuits or lines.)

21.6.4 Shutting Down and Restarting the Network

The network shuts down automatically as part of the system shutdown procedure. If your system is running, you can shut down the network at your local node without destroying any active logical links in one of two ways:

- Shutting down without terminating logical links

Network Considerations

21.6 Managing a Network Node

The following command allows no new logical links; when all existing links are disconnected, the network is turned off:

```
$ RUN SYS$SYSTEM:NCP
NCP> SET EXECUTOR STATE SHUT
NCP> EXIT
```

- Terminating logical links when shutting down

The following command immediately terminates all logical links and stops the network:

```
$ RUN SYS$SYSTEM:NCP
NCP> SET EXECUTOR STATE OFF
NCP> EXIT
```

To start the network if it is not currently active, you must log in to the SYSTEM account or have the privileges listed at the beginning of the STARTNET.COM command procedures.

To start the network manually, invoke the following command:

```
$ @SYS$MANAGER:STARTNET
```

Enable the same command in the site-specific startup procedure so the network starts each time the operating system is booted. To enable the command, use a text editor to delete the exclamation point at the front of the command line in the command procedure.

After enabling the command, the network starts automatically as part of the system startup. Do not start the network again unless you explicitly shut down the network or remove the network startup line from the site-specific startup procedure.

Managing the Local Area Network (LAN) Software

This chapter describes how the LAN software works and the tasks you must perform to manage the LAN software on your system.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Running the LANACP LAN server process	Section 22.3.1
Managing LAN devices	Section 22.5
Managing the LAN device databases	Section 22.6
Managing the LAN node databases	Section 22.7
Migrating from DECnet MOP to LAN MOP	Section 22.8.2
Using CLUSTER_CONFIG_LAN.COM and LAN MOP	Section 22.8.3
Managing the MOP downline load services	Section 22.9
Initiating the MOP console carrier	Section 22.9.8
Requesting MOP trigger boot	Section 22.9.9

This chapter explains the following concepts:

Concept	Section
Local area networks	Section 22.1
LANACP LAN server process	Section 22.3
LANCP utility	Section 22.4
MOP downline load services	Section 22.8

22.1 Understanding Local Area Networks

A local area network (LAN) provides a communications **channel** designed to connect information processing equipment in a limited area such as a room, a building, or a complex of buildings (for example, a campus). Nodes in a LAN can be linked by the following types of data transmission media:

- **Ethernet**—One of the earliest and the most common LANs. Ethernet can refer to either a general LAN application (for example, Ethernet address) or to the specific CSMA/CD (carrier sense multiple access with collision detection) technology that implements the Intel, Xerox, and Digital intercompany Ethernet specifications.

Managing the Local Area Network (LAN) Software

22.1 Understanding Local Area Networks

- FDDI (Fiber Distributed Data Interface)—The current fiber-optic LAN. FDDI is implemented in three ways:
 - As a high-speed backbone connecting mid-speed LANs such as Ethernet
 - As a high-speed LAN connecting workstations or other devices
 - As a high-speed connection between host computers or from host computers to peripheral equipment, such as those found in a data center
- Token Ring—The IEEE 802.5 standard token passing ring. ♦

Alpha

22.1.1 LAN Characteristics

LAN controllers are devices that, along with additional external hardware, implement the Ethernet, FDDI, or Token Ring specifications. A LAN controller and the local system constitute a node. The LAN controller communicates with the local system through the system bus, and with remote systems implementing the Ethernet, FDDI, or Token Ring specifications through the communication medium.¹

Application programs use the LAN driver's QIO interface to perform I/O operations to and from other nodes on the LAN. For detailed information on the QIO interface, see the *OpenVMS I/O User's Reference Manual*.

Table 22-1 provides a brief summary of the differences between the three types of LAN media.

Table 22-1 Characteristics of LAN Media

Characteristic	Ethernet/802.3	FDDI	Token Ring/802.5
Speed	10 Mb/s	100 Mb/s	4 or 16 Mb/s
Max. Frame Size	1518 bytes	4500 bytes	4500 or 17,800 bytes
Max. Stations	1024	500	260
Max. LAN Size	2.8 km	100 km	300 m

22.1.1.1 Ethernet LANs

Ethernet controllers use a single multiaccess channel with CSMA/CD to provide direct access from the processor to the Ethernet. On the Ethernet, all nodes have equal access.

An Ethernet is a cable to which each system or device is connected by a single line. In an office or other area where personal computers and workstations are located, ThinWire Ethernet or unshielded twisted pair cabling is usually used.

Individual systems can either be connected directly to an Ethernet or gain access to an Ethernet by means of a local area interconnect device, such as a DELNI. A DELNI serves as a concentrator, grouping systems together. Many similar devices, such as hubs and repeaters for the various kinds of cable, also provide the connectivity.

¹ The Ethernet specification is described in *The Ethernet—Data Link Layer and Physical Layer Specification*. The FDDI specifications are available from ANSI. The Token Ring specifications are available from IEEE.

Managing the Local Area Network (LAN) Software

22.1 Understanding Local Area Networks

22.1.1.2 FDDI LANs

As implemented by Digital, FDDI uses a dual ring of trees topology. It uses one ring as the primary ring, the other ring as a backup, and the tree configuration for increased flexibility, manageability, and availability.

FDDI controllers use a fiber-optic or twisted-pair cable to provide direct access from the processor to the FDDI Token Ring. Note that FDDI networks and Ethernet networks can be combined to form a single extended LAN. This lets applications running on a system connected to FDDI communicate with applications that run on a system connected to Ethernet.

An FDDI concentrator provides for the attachment of FDDI devices such as VAX and Alpha nodes or FDDI-Ethernet bridges to the FDDI LAN.

22.1.1.3 Token Ring LANs

Alpha

On Alpha systems, Token Ring controllers use either shielded or unshielded twisted pairs of wire to access the ring. Note that it is difficult to connect a Token Ring LAN directly bridged to any other type of LAN. However, routing protocols to other LANs work easily. ♦

22.1.2 LAN Addresses

Nodes on the LAN are identified by unique addresses. A message can be sent to one, several, or all nodes on the LAN simultaneously, depending on the address used.

Upon application, IEEE assigns a block of addresses to a producer of LAN nodes. Thus, every manufacturer has a unique set of addresses to use. Normally, one address out of the assigned block of physical addresses is permanently associated with each controller (usually in read-only memory). This address is known as the hardware address of the controller. Each controller has a unique hardware address.

A LAN address is 48 bits in length. LAN addresses are represented as six pairs of hexadecimal digits (six bytes) separated by hyphens (for example, AA-01-23-45-67-FF). The bytes are displayed from left to right in the order in which they are transmitted; bits within each byte are transmitted from right to left. In this example, byte AA is transmitted first; byte FF is transmitted last.

A LAN address can be a physical address of a single node or a multicast address, depending on the value of the low-order bit of the first byte of the address (this bit is transmitted first). The two types of node addresses are:

- **Physical address**—The unique address of a single node on a LAN. The least significant bit of the first byte of a physical address is 0. (For example, in physical address AA-00-03-00-FC-00, byte AA in binary is 1010 1010, and the value of the low-order bit is 0.)
- **Multicast address**—A multidestination address of one or more nodes on a given LAN. The least significant bit of the first byte of a multicast address is 1. (For example, in the multicast address 0B-22-22-22-22-22, byte 0B in binary is 0000 1011, and the value of the low-order bit is 1.)

Token Ring devices do not support IEEE 802 standard multicast addresses. They do support functional addresses. A functional address is a locally administered group address that has 31 possible values. Each functional address sets one bit in the third through sixth bytes of the address, and bytes 1 and 2 are 03-00 (C0:00 in bit reversed format). To convert a multicast address to a functional address, use the SET DEVICE/MAP command.

Managing the Local Area Network (LAN) Software

22.2 Managing Local Area Networks

22.2 Managing Local Area Networks

The local area network (LAN) software includes two system management tools that work in conjunction with the OpenVMS LAN driver system software:

- Local Area Network Control Program (LANCP)
- LANACP LAN server process

The LAN system management tools:

- Allow you to set LAN parameters to customize your LAN environment.
- Display LAN settings and counters.
- Provide Maintenance Operations Protocol (MOP) downline load support for devices such as terminal servers, x-terminals, and LAN-based printers, and for booting satellites in a VMScluster environment. This MOP support provides an alternative to the traditional method of using either DECnet for OpenVMS or DECnet/OSI software.

Table 22-2 describes the LAN management software and the functionality supported on systems running OpenVMS Alpha and OpenVMS VAX.

Table 22-2 LAN System Management Enhancements

Utility	Description	OpenVMS Support
LAN Auxiliary Control Program (LANACP)	Runs as a server process whose primary function is to provide MOP downline load service. Other services include maintenance of a LAN volatile device database and a LAN volatile node database.	The LANACP utility provides identical functionality on VAX and Alpha systems running OpenVMS Version 6.2.

(continued on next page)

Managing the Local Area Network (LAN) Software

22.2 Managing Local Area Networks

Table 22-2 (Cont.) LAN System Management Enhancements

Utility	Description	OpenVMS Support		
LAN Control Program (LANCP)	<p>Allows you to control LAN software parameters and obtain information from the LAN software. You can use the LANCP utility to:</p> <ul style="list-style-type: none"> Obtain LAN device counters, revision, and configuration information Change the operational parameters of LAN devices on the system Maintain the permanent and volatile LAN device and node databases Update the firmware on LAN devices Control the LANACP LAN server process (including MOP downline load server related functions) Initiate MOP console carrier connections Send MOP trigger boot requests to other nodes 	<p>OpenVMS Alpha Version 6.1 contains the initial implementation of LANCP, which does not include MOP-related functions.</p> <p>OpenVMS Version 6.2 (VAX and Alpha) adds MOP-related functions and extends some of this capability to VAX systems. The following table shows how the LAN utility functions are supported on VAX and Alpha systems:</p>		
		Function	OpenVMS Alpha V6.2	OpenVMS VAX V6.2
		Update firmware?	Yes	No
		Change operational parameters of LAN devices?	Yes	No
		Display LAN device information?	Yes	Limited
		Support MOP functions?	Yes	Yes

22.3 Understanding the LANACP LAN Server Process

You can run the LANACP LAN server process to provide the following services:

- Maintenance of the LAN volatile node database
- Maintenance of the LAN volatile device database
- MOP downline load

The LANCP utility allows you to issue instructions to the LANACP process.

Three principal files are connected with LANACP:

- SYS\$SYSTEM:LANACP.EXE**

This file is the LANACP utility program.

- SYS\$STARTUP:LAN\$STARTUP.COM**

This file starts the LANACP server process.

- SYS\$MANAGER:SYSTARTUP_VMS.COM**

This file contains an entry that can be used to start LANACP automatically at system startup.

Managing the Local Area Network (LAN) Software

22.3 Understanding the LANACP LAN Server Process

In addition, four system logical names are associated with the LANACP LAN server process, which are described in Table 22-3.

Table 22-3 LANACP System Logical Names

Component	Description
LAN\$DLL	Defines the location of downline load files, where the location of the file is not provided in the load request or explicitly defined in the LAN volatile node database. By default, this is defined as SYS\$SYSROOT:[MOM\$SYSTEM].
LAN\$NODE_DATABASE	Defines the location of the LAN permanent node database. By default, this is defined as SYS\$COMMON:[SYSEXE]LAN\$NODE_DATABASE.DAT.
LAN\$DEVICE_DATABASE	Defines the location of the LAN permanent device database. By default, this is defined as SYS\$SPECIFIC:[SYSEXE]LAN\$DEVICE_DATABASE.DAT.
LAN\$ACP	Defines the location of the LANACP LAN server process log file, containing entries describing changes to the LAN permanent device and node databases, and load request and load status information. By default, this is defined as SYS\$MANAGER:LAN\$ACP.LOG.

22.3.1 Running the LANACP LAN Server Process

To start the LANACP LAN server process, type @SYS\$STARTUP:LAN\$STARTUP at the DCL prompt or include this line in the SYS\$MANAGER:SYSTARTUP_VMS.COM command file to start LANACP automatically at system startup.

The following shows the command line as it appears in SYS\$MANAGER:SYSTARTUP_VMS.COM:

```
$!  
$! To start the LANACP LAN server application, remove the comment delimiter  
$! from the command line below.  
$!  
$! @SYS$STARTUP:LAN$STARTUP  
$!
```

22.3.2 Stopping the LANACP LAN Server Process

To stop the LANACP LAN server process, enter the SET ACP/STOP command at the LANCP utility prompt.

22.4 Understanding the LANCP Utility

The LANCP utility allows you to set and show LAN parameters. Section 22.4.1 describes how to invoke the LANCP utility. Table 22-4 describes LAN functions and provides section references to the LANCP commands that help you perform these functions.

Managing the Local Area Network (LAN) Software

22.4 Understanding the LANCP Utility

Table 22-4 Functions of the LANCP Utility

Task	Section
Managing LAN devices	Section 22.5
Managing LAN device databases	Section 22.6
Managing LAN node databases	Section 22.7
Managing the MOP downline load service	Section 22.9
Initiating a MOP console carrier connection	Section 22.9.8
Sending MOP trigger boot requests	Section 22.9.9

22.4.1 Invoking and Exiting LANCP

Table 22-5 describes the ways you can invoke the LANCP utility (SYS\$SYSTEM:LANCP.EXE).

Table 22-5 Invoking the LANCP Utility

Command	Example
Use the RUN command	At the DCL command prompt, enter: \$ RUN SYS\$SYSTEM:LANCP The LANCP utility responds by displaying the LANCP prompt at which you can enter LANCP commands.
Define LANCP as a foreign command	Either at the DCL prompt or in a startup or login command file, enter: \$ LANCP ::= \$SYS\$SYSTEM:LANCP Then, you can enter the command LANCP at the DCL prompt to invoke the utility and enter LANCP commands. When you enter the LANCP command: <ul style="list-style-type: none"> • Without specifying any command qualifiers, the LANCP utility displays the LANCP prompt at which you can enter commands. • With command qualifiers, the LANCP utility terminates after it executes the command and the DCL command prompt is displayed.
Use the MCR command	At the DCL command prompt, enter: \$ MCR LANCP When you enter the MCR LANCP command: <ul style="list-style-type: none"> • Without specifying any command qualifiers, the LANCP utility displays the LANCP prompt at which you can enter commands. • With command qualifiers, the LANCP utility terminates after it executes the command and the DCL command prompt is displayed.

At the LANCP> prompt, you can enter LANCP commands.

For information about the LANCP utility, enter the HELP command at the LANCP> prompt.

To exit from the LANCP utility, enter the EXIT command or press Ctrl/Z at the LANCP> prompt.

Managing the Local Area Network (LAN) Software

22.4 Understanding the LANCP Utility

22.4.2 LANCP Commands

Table 22-6 summarizes the LANCP commands.

Table 22-6 LANCP Commands

Command	Function
@ (Execute Procedure)	Executes a command procedure.
CLEAR DEVICE	Deletes a device from the LAN volatile device database.
CLEAR MOPDLL	Clears MOP downline load counters for all nodes and devices.
CLEAR NODE	Deletes a node from the LAN volatile node database.
CONNECT NODE	Connects to a LAN device, such as a terminal server, that implements a management interface using the MOP console carrier protocol.
DEFINE DEVICE	Enters a device into the LAN permanent device database or modifies an existing entry.
DEFINE NODE	Enters a node into the LAN permanent node database or modifies an existing entry.
EXIT	Stops execution of LANCP and returns control to the DCI command level.
HELP	Provides online help information about the LANCP utility.
LIST DEVICE	Displays information in the LAN permanent device database.
LIST NODE	Displays information in the LAN permanent node database.
PURGE DEVICE	Deletes a device from the LAN permanent device database.
PURGE NODE	Deletes a node from the LAN permanent node database.
SET ACP	Modifies the operation of the LANACP LAN server process.
SET DEVICE (parameters)	Modifies device parameters.
SET DEVICE (volatile device database)	Enters a device into the LAN volatile device database or modifies an existing entry.
SET NODE	Enters a node into the LAN volatile node database or modifies an existing entry.
SHOW CONFIGURATION	Displays a list of LAN devices on the system.
SHOW DEVICE	Displays information in the LAN volatile device database.
SHOW LOG	Displays recent downline load activity.
SHOW MOPDLL	Displays the current state of MOP downline load services.
SHOW NODE	Displays information in the LAN volatile node database.
SPAWN	Creates a subprocess of the current process.

(continued on next page)

Managing the Local Area Network (LAN) Software

22.4 Understanding the LANCP Utility

Table 22–6 (Cont.) LANCP Commands

Command	Function
TRIGGER NODE	Issues a request to reboot to a remote node.
UPDATE DEVICE	Updates firmware image for a device.

For detailed information about LANCP commands and qualifiers, see the *OpenVMS System Management Utilities Reference Manual: A–L*.

22.4.3 LANCP Miscellaneous Functions

Use the SPAWN command to create a subprocess of the current process. The SPAWN command copies the context of the subprocess from the current process. This allows you to exit temporarily from LANCP without having to restart LANCP when you resume.

The syntax for the SPAWN command is as follows:

SPAWN [optional command line]

You can set up the LANCP utility to execute commands from a command file from within LANCP. The LANCP utility recognizes the command file as the file name preceded by @. The default file name extension is .COM.

22.5 Managing LAN Devices

LAN device management consists of displaying device characteristics and setting device parameters. You can use the LANCP utility to set parameters for the types of LAN devices shown in Table 22–7.

Table 22–7 LAN Devices

LAN	Device Examples	Description
Ethernet	DE425, DE434, DE435, DE436, DE500, DECchip 21040	Allow the selection of media type (type of cable connected) and the speed of connection (Ethernet or FastEthernet). Allow full-duplex operation (point-to-point operation between a similar device or between the device and a switch).
FDDI	DEFTA, DEFPA, DEFAA, DEFEA, DEMFA	Allow full-duplex operation.
Token Ring	DETRA, DW300, DW110	Allow the setting of Token Ring parameters and the definition of source routing and functional address mapping.
All	Any	Allow the setting of generic parameters such as the number of receive buffers.

22.5.1 Displaying System Devices

To display the LAN devices on the system, enter the SHOW CONFIGURATION command using the following syntax:

SHOW CONFIGURATION

Managing the Local Area Network (LAN) Software

22.5 Managing LAN Devices

The following example shows the output from a SHOW CONFIGURATION command that was entered on a node that has three LAN devices, two DE435s, and a DETRA:

```
LANCP> SHOW CONFIGURATION
LAN Configuration:
```

Device	Medium	Default LAN Address	Version
EWA0	CSMA/CD	08-00-2B-E4-00-BF	02000023
EWB0	CSMA/CD	08-00-2B-92-A4-0D	02000023
IRA0	Token Ring	00-00-93-58-5D-32	20000223

The version is the device-specific representation of the actual version. In this example, for two devices on the PCI bus, the actual version is in the low byte (2.3 for the DE435 adapters). A device that does not have a readable version is shown as version zero.

Consult your device-specific documentation to correlate the version returned with a particular hardware or firmware implementation of the device.

22.5.2 Displaying Device Parameters

To display information about a LAN device (in the volatile device database), enter the SHOW DEVICE command using the following syntax:

```
SHOW DEVICE device-name [/qualifiers]
```

Table 22-8 provides a brief description of the SHOW DEVICE command qualifiers.

Note

If you do not specify a qualifier, the utility displays the matching devices without additional information.

Table 22-8 SHOW DEVICE Command Qualifiers

Qualifier	Description
/COUNTERS	Displays device counters.
/MAP	Displays the current configuration of the functional address mapping table.
/PARAMETERS	Displays status and related information about the device.
/REVISION	Displays the current firmware revision of the adapter, if available or applicable.
/SR_ENTRY	Displays the contents of the current source routing cache table.

The following are examples of the SHOW DEVICE command:

Managing the Local Area Network (LAN) Software

22.5 Managing LAN Devices

SHOW DEVICE Examples

1. LANCP> SHOW DEVICE/COUNTERS EXA0

Device Counters EXA0:

Value	Counter
259225	Seconds since last zeroed
5890496	Data blocks received
4801439	Multicast blocks received
131074	Receive failure
764348985	Bytes received
543019961	Multicast bytes received
3	Data overrun
1533610	Data blocks sent
115568	Multicast packets transmitted
122578	Blocks sent, multiple collisions
86000	Blocks sent, single collision
189039	Blocks sent, initially deferred
198120720	Bytes sent
13232578	Multicast bytes transmitted
7274529	Send failure
0	Collision detect check failure
0	Unrecognized frame destination
0	System buffer unavailable
0	User buffer unavailable

This command displays counters for Ethernet device EXA0.

2. LANCP> SHOW DEVICE/MAP ICA0

Multicast to Functional Address Mapping ICA0:

Multicast address	Functional Address	Bit-Reversed
09-00-2B-00-00-04	03-00-00-00-02-00	C0:00:00:00:40:00
09-00-2B-00-00-05	03-00-00-00-01-00	C0:00:00:00:80:00
CF-00-00-00-00-00	03-00-00-08-00-00	C0:00:00:10:00:00
AB-00-00-01-00-00	03-00-02-00-00-00	C0:00:40:00:00:00
AB-00-00-02-00-00	03-00-04-00-00-00	C0:00:20:00:00:00
AB-00-00-03-00-00	03-00-08-00-00-00	C0:00:10:00:00:00
09-00-2B-02-00-00	03-00-08-00-00-00	C0:00:10:00:00:00
09-00-2B-02-01-0A	03-00-08-00-00-00	C0:00:10:00:00:00
AB-00-00-04-00-00	03-00-10-00-00-00	C0:00:08:00:00:00
09-00-2B-02-01-0B	03-00-10-00-00-00	C0:00:08:00:00:00
09-00-2B-00-00-07	03-00-20-00-00-00	C0:00:04:00:00:00
09-00-2B-00-00-0F	03-00-40-00-00-00	C0:00:02:00:00:00
09-00-2B-02-01-04	03-00-80-00-00-00	C0:00:01:00:00:00
09-00-2B-02-01-07	03-00-00-02-00-00	C0:00:00:40:00:00
09-00-2B-04-00-00	03-00-00-04-00-00	C0:00:00:20:00:00
09-00-2B-02-01-00	03-00-00-00-08-00	C0:00:00:00:10:00
09-00-2B-02-01-01	03-00-00-00-10-00	C0:00:00:00:08:00
09-00-2B-02-01-02	03-00-00-00-20-00	C0:00:00:00:04:00
03-00-00-00-00-01	03-00-00-00-00-01	C0:00:00:00:00:80
03-00-02-00-00-00	03-00-02-00-00-00	C0:00:40:00:00:00

This command displays mapping information for Token Ring device ICA0.

Managing the Local Area Network (LAN) Software

22.5 Managing LAN Devices

3. LANCIP> SHOW DEVICE/PARAM IRA0

Device Parameters IRA0:

Value	Parameter
Normal	Controller mode
External	Internal loopback mode
00-00-93-58-5D-32	Hardware LAN address
Token Ring	Communication medium
Enabled	Functional address mode.
No	Full duplex enable
No	Full duplex operational
16	Line speed (megabits/second)
16 Mbps	Ring speed
STP	Line media
Enabled	Early token release
Disabled	Monitor contender
200	SR cache entries
2	SR discovery timer
60	SR Aging Timer
Enabled	Source routing
3	Authorized access priority
AA-00-04-00-92-FF	Upstream neighbor
0	Ring number

This command displays status and parameter information for Token Ring device IRA0.

4. LANCIP> SHOW DEVICE/REVISION FXA0

Device revision FXA0: 05140823

This command displays revision information for FDDI device FXA0.

5. LANCIP> SHOW DEVICE/SR_ENTRY ICA0

Source Routing Cache Table ICA0:

LAN address	State	XmtTmo	RcvTmo	StaleTmo	DiscvTmo
AA-00-04-00-92-FF	LOCAL	00000028	00000028	00000245	00000000

This command displays source routing entry information for Token Ring device ICA0.

22.5.3 Setting Device Parameters

All LAN devices are characterized by a collection of parameters. The parameters define the operational characteristics of a LAN device on the medium to which the device is connected.

To set LAN device parameters directly, enter the SET DEVICE command at the LANCIP> prompt. The LANCIP utility issues this command directly to the specified device (without interaction with the LANACP server process).

The syntax for the SET DEVICE command is:

SET DEVICE *device-name* [/qualifiers]

Table 22-9 provides a brief description of the SET DEVICE command qualifiers that apply directly to LAN devices.

Managing the Local Area Network (LAN) Software

22.5 Managing LAN Devices

Table 22–9 SET DEVICE (parameters) Command Qualifiers

Qualifier	Description
/ALL	Sets data for all LAN devices.
/AGING_TIMER= <i>value</i>	Sets the amount of time in seconds to age source routing cache entries before marking them stale.
/CACHE_ENTRIES= <i>value</i>	Sets the number of entries to reserve for caching source routing address entries.
/CONTENDER	Specifies that the device is to participate in the monitor contention process when it joins the ring.
/DISCOVERY_TIMER= <i>value</i>	Sets the number of seconds to wait for a reply from a remote node when performing the source routing route discovery process.
/EARLY	Enables Early Token Release on the device.
/FULL_DUPLEX	Enables full-duplex operation of a device.
/MAP=(MULTICAST_ADDRESS= <i>address</i> , FUNCTIONAL_ADDRESS= <i>address</i>)	Defines a functional address mapping entry.
/MAX_BUFFERS= <i>value</i>	Sets the maximum number of receive buffers to be allocated and used by the LAN driver for the LAN device.
/MEDIA= <i>value</i>	<ul style="list-style-type: none"> For Token Ring devices: Selects the type of cable that is being used to connect the adapter to the Token Ring Media Access Unit (MAU) for devices that do not automatically detect this. For Ethernet devices: Selects the cable connection.
/MIN_BUFFERS= <i>value</i>	Sets the minimum number of receive buffers to be allocated and used by the LAN driver for the LAN device.
/SOURCE_ROUTING	Enables source routing on the Token Ring device.
/SPEED= <i>value</i>	Sets the speed of the LAN, if multiple speeds are supported.
/SR_ENTRY=(LAN_ADDRESS= <i>address</i> , RI= <i>routing-information</i>)	Statically defines a specific source routed route for a specific node.

The following are examples of the SET DEVICE command:

SET DEVICE Examples

1. LANCP> SET DEVICE/CONTENDER/MEDIA=UTP/NOEARLY/SOURCE ICA0

This command enables monitor contention, UTP cable media, and source routing, and disables early token release for Token Ring device ICA0.

2. LANCP> SET DEVICE/MEDIA=TWIST EWBO

This command sets the media type to twisted pair for the second Tulip Ethernet device.

Managing the Local Area Network (LAN) Software

22.5 Managing LAN Devices

3. LANC> SET DEVICE/ALL/MIN_BUFFERS=12

This command sets the number of receive buffers for all LAN devices to be no less than 12.

22.5.4 Updating Device Firmware

LAN devices contain firmware images in EEPROM or FLASH ROM that you can update using the LANC utility. You can update devices such as the DEMNA, DEMFA, DEFAA, DEFTA, DEFEA, and DEFPA.

Note

You can also use methods other than the LANC utility to update firmware. For example, you can use the LFU utility on DEC 7000 and DEC 10000 systems to update DEMNA and DEMFA devices.

To update the firmware on a device, enter the UPDATE DEVICE command using the following syntax:

UPDATE DEVICE *device-name* [/qualifiers]

Table 22-10 provides a brief description of the UPDATE DEVICE command qualifiers.

Table 22-10 UPDATE DEVICE Command Qualifiers

Qualifier	Description
FILE= <i>filename</i>	Provides the file specification of the file to be loaded into the device.
/RESET	Specifies whether the device will begin using the new image when the firmware update completes.

For example, the following command updates FDDI device FAA0 with the firmware image FBUS_MAIN.SYS located on DKA0:[FW]. The device begins using the new image after the firmware update has completed and a device reset has been done.

```
LANC> UPDATE DEVICE FAA0/FILE=DKA0:[FW]FBUS_MAIN.SYS
```

22.6 Managing the LAN Device Databases

The LAN volatile and permanent device databases contain a single entry for each LAN device that exists on the system. Each entry in the LAN volatile device database contains device information and MOP downline load counters information. Each entry in the LAN permanent device database contains device information that is used to populate the volatile database when the LANACP LAN server process is started.

Typically, each database contains the same devices. However, the permanent database may contain entries for devices that have not yet been configured or installed in the system. The LANACP LAN server process maintains the volatile device database. The LANC utility maintains the permanent device database. You can manipulate either database using the LANC utility commands depending on your user privileges, as follows:

Managing the Local Area Network (LAN) Software

22.6 Managing the LAN Device Databases

- Privileged users can add or delete device entries from each database, enable or disable MOP downline load service, and clear MOP downline load counters information for LAN devices
- Unprivileged users can view the MOP downline load status and counters information

The following sections describe how to enter and remove devices from the LAN permanent and volatile device databases, and how to enable and disable MOP downline load services.

22.6.1 Displaying Devices in the LAN Device Databases

To display information in the LAN permanent device database, enter the LIST DEVICE command using the following syntax:

```
LIST DEVICE device-name [/qualifiers]
```

To display information in the LAN volatile device database, enter the SHOW DEVICE command using the following syntax:

```
SHOW DEVICE device-name [/qualifiers]
```

Table 22–11 provides a brief description of the LIST DEVICE and SHOW DEVICE qualifiers.

Table 22–11 LIST DEVICE and SHOW DEVICE Command Qualifiers

Qualifier	Description
/COUNTERS†	Displays device counters.
/MAP†	Displays the current configuration of the functional address mapping table.
/MOPDLL	Displays MOP downline load information.
/PARAMETERS†	Displays status and related information about the device.
/REVISION†	Displays the current firmware revision of the adapter, if available or applicable.
/SR_ENTRY†	Displays the contents of the current source routing cache table.
†SHOW DEVICE only	

Note

If you do not specify a qualifier, the utility displays the matching devices without additional information.

22.6.2 Entering Devices into the LAN Device Databases

To enter a device into the LAN permanent device database or to modify an existing entry, enter the DEFINE DEVICE command using the following syntax:

```
DEFINE DEVICE device-name [/qualifiers]
```

To enter a device into the LAN volatile device database or to modify an existing entry, enter the SET DEVICE command using the following syntax:

```
SET DEVICE device-name [/qualifiers]
```


Managing the Local Area Network (LAN) Software

22.6 Managing the LAN Device Databases

Table 22–12 provides a brief description of the DEFINE DEVICE and SET DEVICE command qualifiers.

Note

Defaults apply to creation of an entry in the device database. If an existing entry is being modified, fields not specified are not changed.

Table 22–12 DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description
/ALL	Defines data for all LAN devices in the LAN permanent or volatile device database.
/MOPDLL=(<i>enable-option</i> , <i>exclusive-option</i> , <i>size-option</i> , <i>knownclientonly-option</i>)	<p>Provides the MOP downline load service settings for the device.</p> <p>In this qualifier, you can specify:</p> <ul style="list-style-type: none">• <i>enable-option</i> Indicates that MOP downline load service should be enabled or disabled for the device.• <i>exclusive-option</i> Indicates that no other provider of MOP downline load service is allowed on the specified LAN device at the same time as LANACP.• <i>knownclientonly-option</i> Indicates that MOP downline load requests should be serviced only for clients defined in the LAN volatile node database.• <i>size-option</i> Specifies the size in bytes of the file data portion of each downline load message.
/UPDATE	Adds LAN devices that are not currently in one of the LAN device databases to that database. The DEFINE DEVICE command applies to the permanent database; the SET DEVICE command applies to the volatile database.
/VOLATILE_DATABASE (DEFINE command only)	Updates the device entries in the LAN permanent device database with any data currently set in the volatile database.
/PERMANENT_DATABASE (SET command only)	Updates the device entries in the LAN volatile device database with any data currently set in the permanent database.

The following examples show how to use the DEFINE DEVICE and SET DEVICE commands:

Managing the Local Area Network (LAN) Software

22.6 Managing the LAN Device Databases

DEFINE DEVICE and SET DEVICE Examples

1. LANCP> DEFINE DEVICE EXA0/MOPDLL=(ENABLE,EXCLUSIVE)

This command defines LAN device EXA0 to enable LANACP MOP downline load service in exclusive mode. The settings of the KNOWNCLIENTSONLY and SIZE characteristics are not changed. If the device entry does not currently exist in the LAN permanent device database, these settings will be set to the defaults.

2. LANCP> DEFINE DEVICE/ALL/MOPDLL=NOEXCLUSIVE

This command sets all LAN devices defined in the LAN permanent device database to nonexclusive mode for LANACP MOP downline load service.

3. LANCP> SET DEVICE EXA0/MOPDLL=(ENABLE,NOEXCLUSIVE)
LANCP> SET DEVICE FXA0/MOPDLL=(ENABLE,EXCL,KNOWN)

These commands enable LANACP MOP downline load service for:

- LAN device EXA0 in nonexclusive mode
- LAN device FXB0 in exclusive mode for only known clients

22.6.3 Deleting Devices from the LAN Device Databases

To delete a device from the LAN permanent device database, enter the PURGE DEVICE command using the following syntax:

PURGE DEVICE *device-name* [/ALL]

To delete a device from the LAN volatile device database, enter the CLEAR DEVICE command using the following syntax:

CLEAR DEVICE *device-name* [/ALL]

For the PURGE DEVICE and CLEAR DEVICE commands, the /ALL qualifier deletes all LAN devices in the LAN permanent device database.

The following examples show how to use the PURGE DEVICE and CLEAR DEVICE commands:

PURGE DEVICE and CLEAR DEVICE Examples

1. LANCP> PURGE DEVICE/ALL

This command deletes all devices from the LAN permanent device database.

2. LANCP> CLEAR DEVICE EXA0

This command deletes device EXA0 from the LAN volatile device database.

22.7 Managing the LAN Node Databases

The LAN volatile and permanent node databases contain a single entry for each defined LAN node. Each entry in the LAN volatile node database contains node information and MOP downline load counters information. Each entry in the LAN permanent node database contains node information that is used to populate the volatile database when the LANACP LAN server process is started.

Managing the Local Area Network (LAN) Software

22.7 Managing the LAN Node Databases

Typically, each database contains the same nodes. The LANACP LAN server process maintains the volatile node database. The LANCP utility maintains the permanent node database. You can manipulate either database using the LANCP utility commands depending on your user privileges, as follows:

- Privileged users can add or delete node entries from each database and clear MOP downline load counters information for LAN nodes
- Unprivileged users can view the node information and MOP downline load status and counters information

The following sections describe how to enter nodes into and remove nodes from the LAN permanent and volatile node databases.

22.7.1 Displaying Nodes in the LAN Node Databases

To display information in the LAN permanent node database, enter the LIST NODE command using the following syntax:

```
LIST NODE device-name [/ALL]
```

To display information in the LAN volatile node database, enter the SHOW NODE command using the following syntax:

```
SHOW NODE device-name [/ALL]
```

For the LIST NODE and SHOW NODE commands, the /ALL qualifier displays data for all nodes in the LAN permanent or volatile node database.

22.7.2 Entering Nodes into the LAN Node Databases

To enter a node into the LAN permanent node database or to modify an existing entry, enter the DEFINE NODE command using the following syntax:

```
DEFINE NODE node-name [/qualifiers]
```

To enter a node into the LAN volatile node database or to modify an existing entry, enter the SET NODE command using the following syntax:

```
SET NODE node-name [/qualifiers]
```

Table 22-13 provides a brief description of the DEFINE NODE and SET NODE command qualifiers.

Table 22-13 DEFINE NODE and SET NODE Command Qualifiers

Qualifier	Description
/ALL	Defines data for all nodes in the LAN permanent or volatile node database.
/ADDRESS= <i>node-address</i>	Associates a LAN address with the node name.
/BOOT_TYPE=VAX_SATELLITE ALPHA_SATELLITE OTHER	Indicates the type of processing required for downline load requests.
/FILE= <i>file-specification</i>	Supplies the file name you want to be provided when the downline load request does not include a file name.
/ROOT= <i>directory-specification</i>	Supplies the directory specification to be associated with the file name.

(continued on next page)

Managing the Local Area Network (LAN) Software

22.7 Managing the LAN Node Databases

Table 22-13 (Cont.) DEFINE NODE and SET NODE Command Qualifiers

Qualifier	Description
/SIZE= <i>value</i>	Specifies the size in bytes of the file data portion of each downline load message.
/V3	Forces the server to respond to only MOP Version 3 boot requests from this node.
/VOLATILE_DATABASE (DEFINE command only)	Updates the node entries in the LAN permanent node database with any data currently set in the volatile database.
/PERMANENT_DATABASE (SET command only)	Updates the node entries in the LAN volatile node database with any data currently set in the permanent database.

The following examples show how to use the DEFINE NODE and SET NODE commands:

DEFINE NODE and SET NODE Examples

1. DEFINE NODE GALAXY/ADDRESS=08-00-2B-11-22-33 -
 /FILE=NISCS_LOAD.EXE -
 /ROOT=\$64\$DIA14:<SYS10.> -
 /BOOT_TYPE=VAX_SATELLITE

This command sets up node GALAXY in the LAN permanent node database for booting as a VAX satellite into a VMScluster system.

The NISCS_LOAD.EXE file is actually located on \$64\$DIA14:<SYS10.SYSCOMMON.SYSLIB>. The <SYSCOMMON.SYSLIB> is supplied by the LANACP LAN server process and is not included in the root definition.

2. DEFINE NODE ZAPNOT/ADDRESS=08-00-2B-11-22-33 -
 /FILE=APB.EXE -
 /ROOT=\$64\$DIA14:<SYS10.> -
 /BOOT_TYPE=ALPHA_SATELLITE

This command sets up node ZAPNOT for booting as an Alpha satellite into a VMScluster system.

The APB.EXE file is actually located on \$64\$DIA14:<SYS10.SYSCOMMON.SYSEXEXE>. Note that the <SYSCOMMON.SYSEXEXE> is supplied by the LANACP LAN server process and is not included in the root definition.

3. SET NODE CALPAL/ADDRESS=08-00-2B-11-22-33 -
 /FILE=APB_061.EXE

This command sets up node CALPAL for booting an InfoServer image. It defines the file that should be loaded when a load request without a file name is received from node CALPAL.

Because the file does not include a directory specification, the logical name LAN\$DLL defines where to locate the file. You could give a directory specification using the file name or by using the /ROOT qualifier.

Note that specifying the file name explicitly in the boot command overrides the file name specified in the node database entry.

Managing the Local Area Network (LAN) Software

22.7 Managing the LAN Node Databases

22.7.3 Deleting Nodes from the LAN Node Databases

To delete a node from the LAN permanent node database, enter the PURGE NODE command using the following syntax:

```
PURGE NODE node-name [/ALL]
```

To delete a node from the LAN volatile node database, enter the CLEAR NODE command using the following syntax:

```
CLEAR NODE node-name [/ALL]
```

For the PURGE NODE and CLEAR NODE commands, the /ALL qualifier deletes all LAN nodes in the LAN permanent or volatile node database.

22.8 Understanding LAN MOP

The collection of utilities and startup command files for LANCP and LANACP provide the necessary functionality for MOP downline load service. These utilities and files load cluster satellites, terminal servers, and systems requiring downline load of special images, such as console update images or system software update images (for InfoServer load).

22.8.1 Coexistence with DECnet MOP

The LAN MOP environment provides functionality that is similar to that provided by DECnet. The result is that a system manager can choose which functionality to use, DECnet MOP or LAN MOP. For VMScluster systems, LAN MOP permits the operation of a VMScluster without the presence of DECnet.

LAN MOP can coexist with DECnet MOP in the following ways:

- Running on different systems
For example, DECnet MOP service is enabled on some of the systems on the LAN, and LAN MOP is enabled on other systems.
- Running on different LAN devices on the same system
For example, DECnet MOP service is enabled on a subset of the available LAN devices on the system, and LAN MOP is enabled on the remainder.
- Running on the same LAN device on the same system but targeting a different set of nodes for service
For example, both DECnet MOP and LAN MOP are enabled, but LAN MOP has limited the nodes to which it will respond. This allows DECnet MOP to respond to the remainder.

22.8.2 Migrating from DECnet MOP to LAN MOP

To migrate to LAN MOP, follow these steps:

1. Decide which nodes are to provide MOP downline load service. These may be the same nodes that currently have service enabled for DECnet.
2. Populate the LAN permanent device database by typing the following command at the DCL prompt:

```
MCR LANCP DEFINE DEVICE/UPDATE
```


Managing the Local Area Network (LAN) Software

22.8 Understanding LAN MOP

3. Populate the LAN permanent node database by entering a node definition for each of the cluster satellite nodes and any other nodes that are similarly defined in the DECnet node database. You can enter this data manually or execute the command procedure SYS\$EXAMPLES:LAN\$POPULATE.COM, following the directions and help provided.
4. Disable service on each of the DECnet circuits where it is currently enabled in the volatile database.
5. Enable service on each LAN device in the LAN permanent device database that you would like to use by typing the following command at the DCL prompt for each device:

MCR LANCP DEFINE DEVICE *device-name*/MOPDLL=ENABLE
6. If high performance is required, select a data size of 1482 bytes and only reduce this if some load requests now fail. Alternatively, set up one system to load those clients that require a small data size and set up a different system to load the other clients.
7. Start the LANACP server process by typing the following command at the DCL prompt:

@SYS\$STARTUP:LAN\$STARTUP

To migrate permanently, follow these steps:

1. Disable service on each of the DECnet circuits in the permanent database.
2. Edit SYS\$MANAGER:SYSTARTUP_VMS.COM to start LANACP at system startup.

To migrate back to DECnet MOP, follow these steps:

1. Stop the LANACP server process by entering the following LANCP command:
SET ACP/STOP
2. Reenable service on each of the DECnet circuits in the permanent and volatile databases.
3. Edit SYS\$MANAGER:SYSTARTUP_VMS.COM to disable startup of LANACP at system startup.

Note

Any nodes that you added while booting with LAN MOP will not have been entered in the DECnet node database as targets for downline load, and they will need to be updated when you return to DECnet MOP.

22.8.3 Using CLUSTER_CONFIG_LAN.COM and LAN MOP

A new cluster management command procedure has been provided to facilitate the use of LANCP for LAN MOP booting of satellites. Called CLUSTER_CONFIG_LAN.COM, it resides in SYS\$MANAGER and is a direct parallel to CLUSTER_CONFIG.COM, which is used by cluster managers to configure and reconfigure a VMScluster system. The two procedures perform the same functions, but CLUSTER_CONFIG.COM uses DECnet MOP for downline load, whereas CLUSTER_CONFIG_LAN.COM uses LAN MOP and does not use DECnet for anything. Therefore, when you add a new node, CLUSTER_

Managing the Local Area Network (LAN) Software

22.8 Understanding LAN MOP

CONFIG_LAN.COM does not ask for the node's DECnet node name and address. Instead, it queries for an SCS node name and an SCS node id number.

For your convenience, you can still run CLUSTER_CONFIG.COM. When you execute CLUSTER_CONFIG.COM, it checks whether LANACP for MOP booting is also running. It also checks to see if DECnet is running. If LANACP is running and DECnet is not, then CLUSTER_CONFIG.COM dispatches to CLUSTER_CONFIG_LAN.COM. If CLUSTER_CONFIG.COM discovers that both LANACP and DECnet are running, it asks the user whether LAN MOP booting is being used, and whether it should call CLUSTER_CONFIG_LAN.COM for the user.

22.8.4 Sample Satellite Load

The following shows how to issue commands to the LANCP utility to enable MOP downline load service and to define node ZAPNOT:

```
set acp/opcom
set device eza0/mopdll=enable
set node ZAPNOT/addr=08-00-2B-33-FB-F2/file=APB.EXE-
    /root=$64$DIA24:<SYS11.>/boot=Alpha
```

The following shows the OPCOM messages displayed when you start up the LANACP LAN server process:

```
***** OPCOM 30-OCT-1994 06:47:35.18 *****
Message from user SYSTEM on GALAXY
LANACP MOP Downline Load Service
Found LAN device EZA0, hardware address 08-00-2B-30-8D-1C

***** OPCOM 30-OCT-1994 06:47:35.25 *****
Message from user SYSTEM on GALAXY
LANACP MOP Downline Load Service
Found LAN device EZB0, hardware address 08-00-2B-30-8D-1D

***** OPCOM 30-OCT-1994 06:47:54.80 *****
Message from user SYSTEM on GALAXY
LANACP MOP V3 Downline Load Service
Volunteered to load request on EZA0 from ZAPNOT
Requested file: $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE

***** OPCOM 30-OCT-1994 06:48:02.38 *****
Message from user SYSTEM on GALAXY
LANACP MOP V3 Downline Load Service
Load succeeded for ZAPNOT on EZA0
System image, $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE (Alpha image)
```

The following display shows the contents of the LAN\$ACP.LOG file:

```
30-OCT-1994 06:47:35.02 Found LAN device EZA0, hardware address 08-00-2B-30-8D-1C
30-OCT-1994 06:47:35.18 Found LAN device EZB0, hardware address 08-00-2B-30-8D-1D
30-OCT-1994 06:47:35.25 LANACP initialization complete
30-OCT-1994 06:47:45.39 Enabled LAN device EZA0 for MOP downline load service in exclusive mode
30-OCT-1994 06:47:54.70 Volunteered to load request on EZA0 from ZAPNOT
    Requested file: $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE
30-OCT-1994 06:48:02.23 Load succeeded for ZAPNOT on EZA0
    MOP V3 format, System image, $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE
    Packets: 2063 sent, 2063 received
    Bytes: 519416 sent, 4126 received, 507038 loaded
    Elapsed time: 00:00:07.42, 68276 bytes/second
```


22.8.5 Cross-Architecture Booting

The LAN enhancements permit cross-architecture booting in a VMScluster system. VAX boot nodes can provide boot service to Alpha satellites, and Alpha boot nodes can provide boot service to VAX satellites. Note that each architecture must include a system disk that is used for installations and upgrades.

22.9 Managing the LAN MOP Downline Load Service

The LANACP LAN server process maintains the LAN volatile node and device databases. The LANCP utility provides commands that:

- Display MOP downline load status and counters information
- Clear counters information
- Enable or disable OPCOM messages and packet tracing

Counters and status information is maintained for each node and device. Counters information includes transmitted and received byte and packet counts, transmit errors, logical errors such as protocol violations and timeouts, and number of load requests. Status includes the time of the last load and the status of the last load.

22.9.1 Enabling MOP Downline Load Service

To enable MOP downline load service, enter the SET DEVICE command using the following syntax:

```
SET DEVICE device-name/MOPDLL=ENABLE
```

In this command, use the *device-name* parameter to supply the LAN controller device name.

See Section 22.6.2 for a complete description of this command.

22.9.2 Disabling MOP Downline Load Service

To disable MOP downline load service, enter the SET DEVICE command using the following syntax:

```
SET DEVICE device-name/MOPDLL=DISABLE
```

In this command, use the *device-name* parameter to supply the LAN controller device name.

See Section 22.6.2 for a complete description of this command.

22.9.3 Displaying the Status and Counters Data

To display MOP downline load status, enter the SHOW MOPDLL command using the following syntax:

```
SHOW MOPDLL
```

The following display shows counters information for a particular node:

LAN MOP DLL Status:

EXA enabled in exclusive mode for known nodes only, data size 1482 bytes
FXA disabled

	#Loads	Packets	Bytes	Last load time	Last loaded
EXA	5	1675	4400620	23-SEP-1994 10:27.51	GALAXY
FXA	0	0	0		

Managing the Local Area Network (LAN) Software

22.9 Managing the LAN MOP Downline Load Service

On this node are two LAN devices, EXA (DEMNA) and FXA (DEMFA). MOP downline load service is enabled on EXA in exclusive mode.

Requests are answered only for nodes that are defined in the LANACP node database. The image data size in the load messages is 1482 bytes. There have been five downline loads, the last one occurring on node GALAXY at 10:27. Finally, no downline loads are recorded for FXA, which is currently disabled for downline load service.

To display recent downline load activity that has been logged in the LAN\$ACP.LOG file, enter the SHOW LOG command using the following syntax:

```
SHOW LOG
```

22.9.4 Displaying the Status and Counters Data for Individual Nodes

To display MOP downline load information for nodes in the LAN permanent node database, enter the LIST NODE command using the following syntax:

```
LIST NODE node-name [/qualifiers]
```

To display MOP downline load status and counters information for nodes in the LAN volatile node database, enter the SHOW NODE command using the following syntax:

```
SHOW NODE node-name [/qualifiers]
```

Table 22-14 provides a brief description of the LIST NODE and SHOW NODE command qualifiers.

Table 22-14 LIST NODE and SHOW NODE Command Qualifiers

Qualifier	Description
/ALL	Displays information for all nodes in the database.
/OUTPUT= <i>command-file-name</i>	Indicates that the output should be directed to the specified file in the form of a list of DEFINE NODE or SET NODE commands. The resulting command file can be used to create the LAN node databases.
/TOTAL (SHOW NODE command only)	Displays counter totals only.

The following example shows output from a command issued on a local node on which there are three nodes defined (GALAXY, ZAPNOT, and CALPAL). CALPAL has issued two load requests:

- The first request is the multicast request from CALPAL that the local node volunteered to accept.
- The second request is the load request sent directly to the local node by CALPAL for the actual load data. The elapsed time from the second load request to completion of the load was 6.65 seconds.

Node Listing:

```
GALAXY (08-00-2B-2C-51-28):  
  MOP DLL: Load file:  APB.EXE  
            Load root:  $64$DIA24:<SYS11.>  
            Boot type:  Alpha satellite
```


Managing the Local Area Network (LAN) Software

22.9 Managing the LAN MOP Downline Load Service

ZAPNOT (08-00-2B-18-7E-33):

MOP DLL: Load file: NISCS_LOAD.EXE
Load root: LAVC\$SYSDEVICE:<SYS10.>
Boot type: VAX satellite

CALPAL (08-00-2B-08-9F-4C):

MOP DLL: Load file: READ_ADDR.SYS
Last file: LAN\$DLL:APB_X5WN.SYS
Boot type: Other
2 loads requested, 1 volunteered
1 succeeded, 0 failed
Last request was for a system image, in MOP V4 format
Last load initiated 30-OCT-1994 09:11:17 on EXA0 for 00:00:06.65
527665 bytes, 4161 packets, 0 transmit failures

Unnamed (00-00-00-00-00-00):

Totals:

Requests received	2
Requests volunteered	1
Successful loads	1
Failed loads	0
Packets sent	2080
Packets received	2081
Bytes sent	523481
Bytes received	4184
Last load	CALPAL at 30-OCT-1994 09:11:17.29

22.9.5 Clearing the Counters Data

To clear MOP downline load counters for all nodes and devices, enter the CLEAR MOPDLL command using the following syntax:

CLEAR MOPDLL

22.9.6 OPCOM Messages

By default, OPCOM messages are enabled. Messages are generated by the LANACP LAN server process when device status changes, load requests are received, and loads complete. These messages are displayed on the operator's console and included in the log file written by LANACP, SYS\$MANAGER:LAN\$ACPL.LOG.

To enable OPCOM messages, enter the SET ACP/OPCOM command using the following syntax:

SET ACP/OPCOM

22.9.7 Load Trace Facility

If the error data produced by the LANACP LAN server process for a load request is not sufficient to help you determine why the load is failing, you can direct the server process to record trace data. The data consists of transmit and receive packet information for every transmit and receive done by the server, and written to a log file for each load attempt. The name of the log file is SYS\$MANAGER:LAN\$nodename.LOG. You can record either all packet data or only the first 32 bytes of each packet.

The following list describes the typical load sequence:

1. Receive a Program Request message on the Load Assistance Multicast Address from the requesting node, code 8.
2. Transmit an Assistance Volunteer message to the requesting node, code 3.

Managing the Local Area Network (LAN) Software

22.9 Managing the LAN MOP Downline Load Service

3. Receive a Program Request message on your node address from the requesting node, code 8.
4. Transmit a Memory Load message to the requesting node with sequence number zero, code 2.
5. Receive a Request Memory Load message requesting the next sequence number (modulo 256), code 10 (decimal).
6. Repeat steps 4 and 5 until there is no more data to send.
7. Transmit a Memory or Parameter Load with Transfer Address message, code 0 or 20 (decimal).
8. Receive a final Request Memory Load message requesting the next sequence number (modulo 256) indicating that the last message has been received, code 10 (decimal).

For cluster satellite loads, the last Memory Load message contains cluster parameters. This message and the final Load with Transfer Address messages are displayed in full even if only partial trace echo has been enabled.

To enable partial tracing of packet data, enter the SET ACP/ECHO command using the following syntax:

```
SET ACP/ECHO
```

To enable full tracing of packet data, add the /FULL qualifier:

```
SET ACP/ECHO/FULL
```

22.9.8 MOP Console Carrier

Console carrier provides a mechanism to connect to a LAN device, such as a terminal server, that implements a management interface using the MOP console carrier protocol. The LANCP utility provides this function in the form of a CONNECT NODE command.

The command syntax is:

```
CONNECT NODE node-specification [/qualifiers]
```

Table 22-15 provides a brief description of the CONNECT NODE command qualifiers.

Table 22-15 CONNECT NODE Command Qualifiers

Qualifier	Description
/DEVICE= <i>device-name</i>	Specifies the LAN controller device name to be used for the connection.
/DISCONNECT= <i>disconnect-character</i>	Specifies a character that you can use to terminate the connection to the remote node.
/PASSWORD=16hexdigits	Supplies the password to be used when the connection is initiated.
/V3 or /V4	Indicates that MOP Version 3 or Version 4 formatted messages, respectively, are to be used to make the connection.

The following examples show how to use the CONNECT NODE command:

Managing the Local Area Network (LAN) Software

22.9 Managing the LAN MOP Downline Load Service

CONNECT NODE Examples

1. `CONNECT NODE GALAXY/DEVICE=EWA0`

This command attempts a console carrier connection to node GALAXY using the Ethernet device EWA0.

2. `CONNECT NODE 08-00-2B-11-22-33/DEVICE=EWA0/PASSWORD=0123456789ABCDEF`

This command attempts a console carrier connection to the given node address using the Ethernet device EWA0, with a password.

22.9.9 MOP Trigger Boot

Some systems recognize and respond to MOP remote boot requests. These systems typically require a password or other mechanism to prevent unwanted boot requests from triggering a reboot of the system. The LANCP utility provides this function in the form of the TRIGGER NODE command.

To request a reboot of a LAN system, enter the TRIGGER NODE command using the following syntax:

`TRIGGER NODE node-specification [/qualifiers]`

Table 22-16 provides a brief description of the TRIGGER NODE command qualifiers.

Table 22-16 TRIGGER NODE Command Qualifiers

Qualifier	Description
<code>/DEVICE=device-name</code>	Specifies the LAN controller device name to be used for sending the boot messages.
<code>/PASSWORD=16hexdigits</code>	Supplies the password to be used when the connection is initiated.

Rather than specify the format to send MOP Version 3 or 4, the LANCP utility sends one message in each format to the target node.

The following examples show how to use the TRIGGER NODE command:

TRIGGER NODE Examples

1. `TRIGGER NODE GALAXY/DEVICE=EWA0`

This command sends MOP trigger boot messages to node GALAXY using Ethernet device EWA0.

2. `TRIGGER NODE 08-00-2B-11-22-33/DEVICE=EWA0/PASSWORD=0123456789ABCDEF`

This command sends MOP trigger boot messages to the given node address using the Ethernet device EWA0, with indicated password.

THE JOURNAL OF THE ROYAL SOCIETY OF MEDICINE

Volume 68, No. 1, February 1975

Published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1

Subscription prices: £12.00 per annum in advance (UK); £14.00 (overseas)

Single issues: 50p (UK); 60p (overseas)

255 THE JOURNAL OF THE ROYAL SOCIETY OF MEDICINE

The Journal of the Royal Society of Medicine is a peer-reviewed journal of medicine and surgery. It is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

THE JOURNAL OF THE ROYAL SOCIETY OF MEDICINE

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

THE JOURNAL OF THE ROYAL SOCIETY OF MEDICINE

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

THE JOURNAL OF THE ROYAL SOCIETY OF MEDICINE

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1. The journal is published quarterly, in February, May, August, and November. The journal is published by the Royal Society of Medicine, 11, St Andrews Place, Regents Park, London, N.W.1.

Managing InfoServer Systems

This chapter describes InfoServer functions and InfoServer Client for OpenVMS software, which enables OpenVMS systems to access InfoServer device services. The chapter also describes the tasks you must perform to start the client software on your system and to make InfoServer devices available as public devices.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Establishing a server management session	Section 23.3
Starting InfoServer Client for OpenVMS software automatically	Section 23.5.3
Making InfoServer devices available automatically	Section 23.6.3

This chapter explains the following concepts:

Concept	Section
InfoServer functions	Section 23.1
LASTport protocols	Section 23.2
InfoServer Client for OpenVMS functions	Section 23.4
LASTCP utility functions	Section 23.5
LADCP utility functions	Section 23.6

23.1 Understanding InfoServer Functions

The InfoServer system is a high-performance **virtual device server**. It can make available, or **serve**, compact discs, read/write disks, magneto-optical (MO) devices, and tapes to client systems on the local area network (LAN). Systems running InfoServer Client software can connect to the virtual devices and use them as though they are locally attached devices.

Unlike a **file server**, the InfoServer system does not impose a file system on the virtual devices that it serves. For example, the InfoServer system can serve a disk with any type of on-disk file structure. The client system interprets the on-disk structure and uses its own native file system to access data. Multiple on-disk structures can be served by and accessed on a single InfoServer system at the same time.

Managing InfoServer Systems

23.1 Understanding InfoServer Functions

The InfoServer system can perform the following functions:

- Serve compact discs

The InfoServer system serves compact discs automatically, using a disc's volume label as the service name when the server is booted or when a disc is inserted into an InfoServer drive. You do not have to perform any management action. Client systems simply bind to and mount the disc under its volume label.

The InfoServer system can automatically serve to OpenVMS clients compact discs that are in ODS-2 format. High Sierra and ISO-9660 compact discs and other media types can be served manually through the InfoServer management interface.

- Serve Small Computer System Interface (SCSI) tapes

Using service names, the InfoServer system can serve SCSI tape devices to the network. Client systems can connect to these tape devices and use them as though they were locally attached devices.

- Serve read/write disk partitions

A **partition** is a logical subset of an InfoServer read/write disk. A single disk can be subdivided into several partitions, each of which can be served to the network independently. To remote client systems, these partitions appear to be whole disks. For example, a client system using InfoServer Client for OpenVMS software can access the partitions and use them as though they are local hard disks.

- Act as an initial load system for OpenVMS systems

The InfoServer system can downline load the primary bootstrap program to OpenVMS systems by responding to maintenance operation protocol (MOP) requests. The server can locate MOP downline load files on the OpenVMS software distribution compact disc and copy them into temporary MOP partitions on an InfoServer-formatted read/write disk.

The initial system load (ISL) bootstrap program connects back to the software distribution compact disc and boots Standalone Backup. The Backup utility is then used to copy the OpenVMS operating system save sets from the compact disc to a read/write disk attached to the system. All subsequent OpenVMS boots are done from the local read/write disk.

- Downline load other products

You can use the InfoServer system to load any Ethernet product by file name; that is, the server does not require a Network Control Program (NCP) database entry to locate the requested file. For example, X terminal clients use the InfoServer system to downline load their system software. You can create a special MOP partition and copy the desired file to that partition. The server additionally supports downline loading of services by Ethernet address. Each InfoServer system can handle up to 100 simultaneous downline loads more efficiently than host-based downline loaders, which must start processes to assist in the load.


Figure 23-1 shows the relationship of the InfoServer system to several possible client systems. In this figure, two compact discs and two hard disks connected to the server appear to the client systems as local devices. The VAX system and the RISC workstation might be using one or two of the compact discs for software distribution and online documentation, while the PC might be referencing a disk

Managing InfoServer Systems

23.1 Understanding InfoServer Functions

partition on the InfoServer system. The X terminal boots from the InfoServer system and uses InfoServer disks for page, font, and customization files.

Figure 23-1 InfoServer System Serving Clients



ZK-5901A-GE

You can connect the InfoServer system to your Ethernet LAN and turn on the system. After the server is initialized, or **bootstrapped**, the server software automatically serves to client systems the device media connected to it. If you insert a compact disc into a server drive, the server detects this new device and automatically serves it to client systems by using the volume label as the service name.

The server bootstraps from its internal read/write device, on which the InfoServer software is preinstalled. InfoServer software updates are distributed on compact discs. As these new releases become available, you can install the software onto the internal device for subsequent booting. To update InfoServer software from the compact disc, follow these steps:

1. Insert the disc in a compact disc drive attached to the InfoServer system.
2. Move the InfoServer software to the internal read/write device. At the InfoServer prompt, enter a command in the following format, where *n* is the drive number:

On the InfoServer 100 or InfoServer 150 system:

```
InfoServer> UPDATE SYSTEM DKn:
```

On the InfoServer 1000 system:

```
InfoServer> UPDATE SYSTEM DKn: FLASH
```

The next time you boot the InfoServer system, it runs the updated software.

You might want to customize server features. You can control InfoServer functions by logging in to the server and entering server commands, described in detail in the *InfoServer System Operations Guide*.

Managing InfoServer Systems

23.1 Understanding InfoServer Functions

23.1.1 Automatic Service Policies for Multiple Servers

The InfoServer system automatically serves its locally connected devices to clients when the server is first powered on or when a removable device (for example, a compact disc) is inserted into a drive. The server reads the volume label of each device and uses the label as the name of the service offered to clients.

Note

You can disable the automatic service feature by using the InfoServer command SET SERVER AUTOMOUNT.

If multiple servers offer the same services, the client uses a rating scheme to select the appropriate service. See the CREATE SERVICE command description in the *InfoServer System Operations Guide* for more information.

When you remove a compact disc from a server drive, the InfoServer system ends all client connections to the associated service. The InfoServer system also stops offering the associated service to client systems.

23.1.2 High-Availability Feature to Reduce Service Interruptions

The InfoServer system provides a high-availability feature that is especially beneficial for OpenVMS clients. If the server ends a service connection for some reason (for example, the server reboots, or you remove a compact disc), the OpenVMS client enters mount verification for that volume. If the same service is offered by another InfoServer system on the LAN, the client automatically connects to that service.

For example, suppose you have two identical copies of the OpenVMS Online Documentation compact disc in drives on two different servers. If one server or drive fails, a new connection is established to the duplicate disc on the other server. File operations continue as normal, and users experience almost no service disruption.

23.1.3 Support for X Terminal Clients

X terminal clients use the InfoServer system to download their system software, provide font services, save configuration information, and page memory to and from InfoServer disks. For example, system files for Digital's VXT 2000 windowing terminals can be installed from compact disc on the InfoServer system. Once installed, these files are downline loaded on demand to each terminal when it is powered on.

The terminals can dynamically allocate partitions on an InfoServer disk as needed. For example, when a user requests that terminal customizations be saved, the InfoServer system automatically creates a disk partition to hold the information and creates a network service name for that partition. Once customization information is saved, the user can recall the information at any time.

VXT 2000 terminals that are InfoServer clients can also be virtual memory machines. Such terminals can page sections of main memory to and from InfoServer disks as required. Because a VXT client has no local disk, it uses InfoServer disks as page disks. When main memory is paged out to disk, the VXT client requests the InfoServer system to create a partition. This partition is then automatically extended as needed. Partitions and their network service names are created dynamically, without requiring user action.

Managing InfoServer Systems

23.1 Understanding InfoServer Functions

By default, the InfoServer disk DK1, which is the internal disk that ships with each InfoServer 150 system, is enabled to allow VXT 2000 clients to allocate partitions remotely. Other disks can also be enabled through the use of InfoServer commands.

23.2 Understanding LASTport Protocols

The InfoServer system uses the LASTport transport protocol and the LASTport/Disk and LASTport/Tape system application protocols to provide access to the virtual devices it serves to the LAN. These protocols provide high-performance access to disk and tape devices. The InfoServer system implements the server portion of the protocols, while the client systems that access InfoServer storage devices implement the client portion.

On OpenVMS systems running the LASTport transport, all Ethernet devices must be terminated either by attaching the devices to an active network or by using an appropriate terminator. Failure to terminate the devices causes a system crash.

23.2.1 LASTport Transport Protocol

The LASTport protocol is a specialized LAN transport protocol that allows many clients to access InfoServer systems and perform reliable transactions. For the InfoServer system, a transaction is a device read or write operation. The LASTport protocol allows many client systems concurrently to read information from, and write information to, an InfoServer storage device.

Unlike timer-based protocols, the LASTport protocol is a transaction-oriented protocol. Normally, information does not pass between a client and an InfoServer system unless the client initiates a transaction. The client system then runs a timer on the transaction, normally waiting from two to five seconds before assuming that the transaction is lost and retrying the operation.

The LASTport protocol does not provide any routing functions; it runs only in a LAN. The LASTport protocol type is 80-41. If the extended LAN uses any filtering devices, they must allow this protocol type to pass unfiltered so that clients can access InfoServer systems across the filtering device.

The InfoServer system uses a multicast address feature of the LASTport protocol to establish connections to devices. The format of the multicast address is 09-00-2B-04-*nn-nn*, where *nn* depends on the work group enabled (see the *InfoServer System Operations Guide*).

23.2.2 LASTport/Disk Protocol

The LASTport/Disk protocol is a specialized device protocol that uses the LASTport transport. That is, LASTport/Disk messages are delivered in LASTport messages. The LASTport/Disk protocol provides the mechanism for reading and writing logical blocks independent from any underlying file system. The clients that implement the LASTport/Disk protocol interpret the file system locally. By using the LASTport/Disk protocol for access to compact discs and read/write disks, the InfoServer system can support multiple client operating systems and on-disk structures concurrently.

The LASTport/Disk protocol also provides the naming facility to access compact discs and read/write disks. The InfoServer system assigns each virtual device a service name and allows clients to query the LAN for these names. When the requested service is found, the client connects to it, and device access can begin.

Managing InfoServer Systems

23.2 Understanding LASTport Protocols

When duplicate virtual devices are available under identical service names, the protocol provides a facility for load balancing among the available devices.

23.2.3 LASTport/Tape Protocol

Like the LASTport/Disk protocol, the LASTport/Tape protocol uses the LASTport transport. That is, LASTport/Tape messages are delivered in LASTport messages. The LASTport/Tape protocol provides the mechanism for reading and writing tape records. Tape devices attached to the InfoServer system appear to tape clients as locally attached devices.

The LASTport/Tape protocol also provides the naming facility to access tapes. The InfoServer system assigns each tape device a service name and allows clients to query the LAN for these names. When the requested service is found, the client connects to it, and tape access can begin.

23.3 Establishing a Server Management Session

You can establish a server management session from a local or remote console terminal:

- **For a local session**, you connect a terminal capable of interpreting VT100 ANSI escape sequences to the serial port on the rear of the InfoServer system unit (MMJ1 on the InfoServer 150 unit). The terminal must be set to 9600 baud, 8 bits, no parity.
- **For a remote session**, you make a connection to the InfoServer system through a local area terminal (LAT) server.

Like many network servers, the InfoServer system advertises a LAT service for its management interface and accepts connections from remote terminals attached to terminal servers. Therefore, any terminal attached to a terminal server on the extended LAN can act as a console terminal for the InfoServer system (if the user knows the InfoServer management password).

Determining the Server's Default Service Name

To make a remote connection to the InfoServer system for the first time, you must determine the server's default name. To do this, add the four-character prefix LAD_ to the hexadecimal Ethernet datalink address on the InfoServer system's cabinet. You can change this default name by using the InfoServer command SET SERVER NAME.

The server's name is the LAT service name to which you connect. For example, if the default server name is LAD_08002B15009F, you would enter the following command at the terminal server's prompt to manage the InfoServer system:

```
Local> CONNECT LAD_08002B15009F
```

See your terminal server user's guide for information about the establishment of LAT service connections.

Entering an InfoServer Password

After you connect to the InfoServer system, you must enter an InfoServer password to establish the management session. The default server password is ESS. You can change the password with the InfoServer command SET SERVER PASSWORD.

Managing InfoServer Systems

23.3 Establishing a Server Management Session

Example

The following example shows the establishment of a sample session using a DECserver 500 terminal server:

```
Local> CONNECT LAD_08002B133C1C
Password: ESS (not echoed)
Local -010- Session 1 to LAD_08002B133C1C established
```

```
DEC InfoServer V3.1
```

```
InfoServer> SHOW SERVER
```

In this example, the terminal server's prompt is Local>, and a LAT session is established to the InfoServer system whose service name is LAD_08002B133C1C. The InfoServer system prompts for a server password. When you enter the correct password, the server prompts for InfoServer commands with the InfoServer> prompt.

Ending a Session

At the end of the management session, you can enter the EXIT command at the InfoServer> prompt. This command returns you to the terminal server's Local> prompt if the management session is over a LAT connection.

23.3.1 Server Management Commands

Table 23-1 summarizes InfoServer commands and their functions.

Table 23-1 Summary of InfoServer Commands

Command	Function
BACKUP	Saves InfoServer-formatted disks.
BIND	Establishes a connection to the specified ODS-2 service and creates the virtual device VDK1 for that service.
CLEAR	Erases the console terminal screen.
COPY	Copies data from one disk or partition to another.
CRASH	Causes the server software to take a recognizable bugcheck, creating a dump if crashdump processing is enabled.
CREATE	Creates a new partition or service.
DELETE	Deletes a partition or service that was previously created.
DISCONNECT	Terminates a LASTport or LAT terminal server session.
ERASE	Erases the specified disk or partition; erases FUNCTIONS or SERVICES data from non-volatile random-access memory (NVRAM).
EXIT	Terminates the management session.
HELP	Displays help text for the InfoServer commands.
INITIALIZE	Formats a read/write disk into an InfoServer disk.
LOOP	Automatically repeats any valid InfoServer command.
MONITOR	Automatically repeats valid InfoServer commands every 3 seconds, clearing the screen and placing the cursor at the home position.
PURGE	Purges old versions of VXT software.

(continued on next page)

Table 23–1 (Cont.) Summary of InfoServer Commands

Command	Function
REBOOT	Shuts down and reboots the server.
RECORD	Records data from an InfoServer disk or partition to a compact disc.
RESTORE	Resets the server to a previously saved system configuration.
RETRIEVE	Restores InfoServer-formatted disks saved by the BACKUP command.
REWIND	Rewinds an InfoServer tape.
SAVE	Saves configuration and service data for recovery after a server reboot.
SET	Sets partition, service, or server parameters.
SHOW	Displays the server's parameters and counters.
UNBIND	Deletes the VDK1 virtual device and terminates the connection to the remote service.
UNLOAD	Rewinds and unloads an InfoServer tape.
UPDATE	Installs one or more new software products or functions.
VERIFY	Validates the on-disk structure of a device formatted with the INITIALIZE command.
ZERO	Sets internal server counters to 0.

The InfoServer system provides a Help facility that contains information about each server command, including parameters, qualifiers, and examples of its use. For detailed information on InfoServer commands, refer to the *InfoServer System Operations Guide*.

23.4 Understanding InfoServer Client for OpenVMS Functions

InfoServer Client for OpenVMS software enables clients running the OpenVMS operating system to access virtual device services offered by InfoServer systems on a LAN. Software components include the following:

- **LASTport driver**
The LASTport driver provides reliable data transfer services for its clients. It interacts with the Data Link driver and the LASTport/Disk driver as an efficient transport for a virtual device service. The LASTport driver can support other applications, such as a primitive data queueing service.
- **LASTport/Disk client driver**
The LASTport/Disk client driver presents a standard block device interface to the system. The OpenVMS file system interacts with the LASTport/Disk client as if the LASTport/Disk client were a local disk driver. The LASTport/Disk client driver supports both raw and buffered interfaces.
- **LASTport/Tape client driver**
The LASTport/Tape client driver enables OpenVMS clients to access and use as local devices SCSI tapes attached to InfoServer systems.
- **LASTCP and LADCP utilities**

23.4 Understanding InfoServer Client for OpenVMS Functions

These utilities allow you to start InfoServer Client software on your system, monitor transport status, and configure and maintain InfoServer device services. Section 23.5 and Section 23.6 introduce the utilities. For complete information about the utilities, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

23.5 Understanding LASTCP Utility Functions

InfoServer Client for OpenVMS software uses the LASTport protocol to communicate with InfoServer systems on an extended LAN. The protocol is implemented in the OpenVMS device driver ESS\$LASTDRIVER.

The LASTport Control Program (LASTCP) utility is the management interface that allows you to control and diagnose ESS\$LASTDRIVER. You can use LASTCP to do the following:

- Start and stop ESS\$LASTDRIVER
- Display counters for circuits, lines, nodes, and ESS\$LASTDRIVER
- Display node characteristics
- Display known clients and servers
- Display LASTport status
- Reset counters

The description of the LASTCP utility covers the following topics:

- Invoking and exiting the utility
- LASTCP command summary
- Starting InfoServer Client for OpenVMS software automatically

23.5.1 Invoking and Exiting the LASTCP Utility

Use of LASTCP requires normal privileges, except where noted. To invoke LASTCP, enter the following command:

```
$ RUN SYS$SYSTEM:ESS$LASTCP
%LASTCP-I-VERSION, ESS$LASTDRIVER V1.5 is running
LASTCP>
```

At the LASTCP> prompt, you can enter LASTCP commands. To exit the utility, enter EXIT or press Ctrl/z at the LASTCP> prompt.

You can also execute a single LASTCP command by using a DCL string assignment statement, as shown in the following example:

```
$ LASTCP ::= $ESS$LASTCP
$ LASTCP SHOW CLIENTS
```

LASTCP executes the SHOW CLIENTS command and returns control to DCL command level.

Managing InfoServer Systems

23.5 Understanding LASTCP Utility Functions

23.5.2 LASTCP Command Summary

Table 23–2 summarizes LASTCP commands and their functions.

Table 23–2 Summary of LASTCP Commands

Command	Function
EXIT	Returns the user to DCL command level
HELP	Displays HELP text for LASTCP commands
SHOW CIRCUIT COUNTERS	Displays circuit counters
SHOW CLIENTS	Displays known clients
SHOW LINE COUNTERS	Displays line counters
SHOW NODE CHARACTERISTICS	Displays node characteristics
SHOW NODE COUNTERS	Displays node counters
SHOW SERVERS	Displays known servers
SHOW STATUS	Displays local status
SHOW TRANSPORT COUNTERS	Displays transport counters
START TRANSPORT	Starts LASTDRIVER
STOP TRANSPORT	Stops LASTDRIVER
ZERO COUNTERS	Resets counters

You can abbreviate LASTCP commands to the first unique characters of the command verb. For example, you can abbreviate the command SHOW SERVERS to SH SE.

LASTCP provides a Help facility that contains information about each command and its parameters and qualifiers, as well as examples of its use. For a complete description of LASTCP commands, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

23.5.3 Starting InfoServer Client for OpenVMS Software Automatically

You must start InfoServer Client for OpenVMS software using the ESS\$STARTUP command procedure. To make sure the software is started automatically each time the system reboots, execute the startup procedure from within SYSTARTUP_VMS.COM.

How to Perform This Task

1. Determine the value of SCSNODE, your system's node name parameter. If the parameter is defined as the null string (the default value), InfoServer Client for OpenVMS software does not start.

If you are running or plan to run DECnet for OpenVMS, SCSNODE must be defined as the system's DECnet node name. If you do not plan to run DECnet, and if the system is a VMScluster member, SCSNODE must be defined as the SCS system name, a 1- to 8-character node name that is unique in the cluster.

To determine the value of SCSNODE, enter the following commands to invoke SYSMAN and display the parameter:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE CURRENT
SYSMAN> PARAMETERS SHOW SCSNODE
```


Managing InfoServer Systems

23.5 Understanding LASTCP Utility Functions

2. If SCSNODE is defined as the null string, perform these steps:
 - a. Enter a command in the following format, where *node-name* is the system's DECnet node name or (if you do not plan to run DECnet for OpenVMS) the SCS system name:

```
PARAMETERS SET SCSNODE "node-name"
```

For example:

```
SYSMAN> PARAMETERS SET SCSNODE "MYNODE"
```
 - b. Enter the following commands to write the new value to the parameter file and exit from SYSMAN:

```
SYSMAN> PARAMETERS WRITE CURRENT  
SYSMAN> EXIT
```
 - c. Add a line in the following format to the AUTOGEN parameter file `SYS$SYSTEM:MODPARAMS.DAT` to define the SCSNODE parameter:

```
SCSNODE = "node-name"
```

For example:

```
SCSNODE = "MYNODE"
```
3. Invoke any editor to edit `SYS$MANAGER:SYSTARTUP_VMS.COM` and find the command that starts InfoServer Client software. For example:

```
$ @SYS$STARTUP:ESS$STARTUP DISK
```

Note that the parameters `CLIENT` and `DISK` are synonymous. If the command is preceded by the DCL comment delimiter (!), remove the delimiter. If you want to enable tape functions, add the `TAPE` parameter to the command line:

```
$ @SYS$STARTUP:ESS$STARTUP DISK TAPE
```
4. If `SYSTARTUP_VMS.COM` invokes the DECnet for OpenVMS startup procedure (`SYS$MANAGER:STARTNET.COM`), make sure `SYSTARTUP_VMS.COM` executes the InfoServer Client for OpenVMS startup procedure *after* invoking `STARTNET.COM`.

The following example shows the network startup command line followed by the InfoServer Client for OpenVMS startup command line. Note that if you omit the `TAPE` parameter, only the disk function is started.

```
$ @SYS$MANAGER:STARTNET  
:  
:  
:  
$ @SYS$STARTUP:ESS$STARTUP DISK TAPE
```
5. Optionally, edit the file `SYS$STARTUP:ESS$LAST_STARTUP.DAT` to specify desired startup qualifiers for the LASTport transport. (See the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.)

Managing InfoServer Systems

23.5 Understanding LASTCP Utility Functions

23.5.4 Startup Restrictions: PATHWORKS and RSM

If PATHWORKS or Remote System Manager (RSM) or both are installed, the InfoServer Client for OpenVMS startup must be run before the startup for PATHWORKS or RSM, or both. For example:

```
$ @SYS$MANAGER:STARTNET
.
.
$ @SYS$STARTUP:ESS$STARTUP DISK TAPE
$ @SYS$STARTUP:PCFS_STARTUP
$ @SYS$STARTUP:RSM$SERVER_STARTUP
```

InfoServer Client for OpenVMS software provides device drivers and control programs that are shared by both the PATHWORKS and RSM products. All InfoServer Client for OpenVMS components are prefixed with ESS\$. The drivers and control programs supplied with InfoServer Client for OpenVMS software provide all necessary support for both PATHWORKS and RSM in addition to InfoServer Client support. You must execute the InfoServer Client for OpenVMS startup in the site-specific startup before executing either the PATHWORKS or RSM startup procedure.

23.5.5 Startup Restrictions: SYSMAN

You cannot start InfoServer Client for OpenVMS from a subprocess. Because the OpenVMS System Management utility (SYSMAN) uses subprocesses to complete its tasks on remote nodes, SYSMAN cannot be used to execute the SYS\$STARTUP:ESS\$STARTUP procedure.

23.5.6 User Account Requirements

To work with InfoServer Client for OpenVMS software, user accounts on your system must have the following privileges and quotas:

- GRPNAM privilege to use the /GROUP qualifier of the LADCP command BIND; SYSNAM privilege to use the command's /SYSTEM qualifier.
- At a minimum, default UAF account quotas.

See the AUTHORIZE section in the *OpenVMS System Management Utilities Reference Manual* for an explanation of how to verify and change account privileges and quotas.

23.5.7 System Parameter MAXBUF Requirement

To use all the LASTport Control Program (LASTCP) utility's SHOW functions, you must set the value of the system parameter MAXBUF to 32000 or greater.

23.6 Understanding LADCP Utility Functions

You use the LAD Control Program (LADCP) utility to configure and control the LASTport/Disk and LASTport/Tape protocols on OpenVMS systems. OpenVMS systems that use LASTport/Disk and LASTport/Tape services are called client systems. You can use LADCP to do the following:

- Establish **bindings** to services. A binding creates a new DAD n : virtual disk unit or a new MAD n : virtual tape unit on the local OpenVMS system.
- Remove bindings to services.

Managing InfoServer Systems

23.6 Understanding LADCP Utility Functions

You can control service access by using a service access password. You can also write-protect services. In this case, local OpenVMS users of a DADn: or MADn: device unit receive an error if they attempt a write operation to the unit.

The protocols allow you to access storage devices that reside on an InfoServer system as though they are locally connected to your OpenVMS system. Thus, several OpenVMS client systems can share the same read-only media, eliminating the need for duplicate drives and media.

DADn: and MADn: device units are also referred to as **virtual device units**. They represent the local OpenVMS context for a volume that resides on a remote server. The OpenVMS driver that controls the DADn: units is called ESS\$DADDRIVER. The OpenVMS driver that controls the MADn: units is called ESS\$MADDRIVER.

The LASTport/Disk and LASTport/Tape protocols depend on the LASTport transport. The ESS\$STARTUP.COM command procedure in SYS\$STARTUP automatically loads ESS\$DADDRIVER and ESS\$MADDRIVER as well as ESS\$LASTDRIVER, the LASTport transport driver.

Note

Your site-specific startup command procedure must include a call to ESS\$STARTUP.COM. If you are using DECnet software, you must place the call *after* the @SYS\$MANAGER:STARTNET.COM command that starts DECnet software. See Section 23.5.3.

23.6.1 Invoking and Exiting the LADCP Utility

To invoke LADCP, enter the following command:

```
$ RUN SYS$SYSTEM:ESS$LADCP
LADCP>
```

You can enter LADCP commands at the LADCP> prompt.

You can also execute a single LADCP command by using a DCL string assignment statement, as shown in the following example:

```
$ LADCP ::= $ESS$LADCP
$ LADCP BIND CD_DOC_00661 /NOWRITE
```

LADCP executes the BIND command and returns control to DCL command level.

To exit LADCP, enter EXIT or press Ctrl/z after the LADCP> prompt.

23.6.2 LADCP Command Summary

Table 23-3 summarizes LADCP commands and their functions.

Table 23-3 Summary of LADCP Commands

Command	Function
BIND	Establishes a service binding and creates a device unit
DEALLOCATE	Terminates any active connection to a service without deleting the unit control block (UCB)

(continued on next page)

Managing InfoServer Systems

23.6 Understanding LADCP Utility Functions

Table 23-3 (Cont.) Summary of LADCP Commands

Command	Function
EXIT	Returns the user to DCL command level
HELP	Displays help text for LADCP commands
SHOW SERVICES	Displays services offered by InfoServer systems on the LAN
UNBIND	Terminates an established service binding

LADCP provides a Help facility that contains information about each LADCP command, including parameters, qualifiers, and examples of its use. For detailed descriptions of LADCP commands, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

23.6.3 Making InfoServer Devices Available Automatically

You can make remote InfoServer devices available on your system each time the system boots. To do so, add to SYSTARTUP_VMS.COM a series of LADCP BIND commands. For more information about the BIND command, see the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

How to Perform This Task

1. Edit SYSTARTUP_VMS.COM and find the command that starts InfoServer Client software. For example:

```
@SYS$STARTUP:ESS$STARTUP DISK TAPE
```

This command starts the software with disk and tape functions.

2. Add the following command to invoke LADCP:

```
$ RUN SYS$SYSTEM:ESS$LADCP
```

3. Immediately after this command, add BIND commands in the following format to make InfoServer compact discs or read/write disks available as virtual device units:

```
BIND [/QUALIFIER,...] service-name
```

To make tape devices available, you must specify the /TAPE qualifier in addition to any other desired qualifiers:

```
BIND/TAPE [/QUALIFIER,...] service-name
```

For *service-name*, specify the name of the InfoServer device service. Usually a service name is the label of the volume to which the InfoServer system is providing access. For more information on the BIND command, see the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

4. Add an EXIT command to exit LADCP.
5. Add MOUNT commands in the following format to make available as public devices the virtual device units created in the previous step:

```
MOUNT/SYSTEM/NOASSIST device-name volume-label
```

For *device-name*, specify the name of the device. For *volume-label*, specify a volume label to assign to the device. For more information on the MOUNT command, see the MOUNT section in the *OpenVMS System Management Utilities Reference Manual*.

Managing InfoServer Systems

23.6 Understanding LADCP Utility Functions

Example

The following commands, executed in SYSTARTUP_VMS.COM, start the InfoServer Client software and make available to client systems the InfoServer device DAD\$VMS055.

```
.  
. .  
$ @SYS$STARTUP:ESS$STARTUP DISK  
$ RUN SYS$SYSTEM:ESS$LADCP  
  BIND VMS055  
  EXIT  
$ MOUNT/SYSTEM/NOASSIST DAD$VMS055 VMS055  
. .  
.
```

In this example, the VMS Version 5.5 consolidated distribution (CONdist) compact disc loaded in a compact disc drive connected to an InfoServer system, is made available on the server as a virtual device unit and mounted as a public device.

Abstract

The purpose of this study was to determine the effect of the use of a 100% cotton fishing net on the catch of the common carp (*Cyprinus carpio*) in the Tigris River, Iraq. The study was conducted during the summer months of 2005 and 2006. The results showed that the use of a 100% cotton fishing net resulted in a higher catch of common carp compared to the use of a mixed cotton and nylon net.

Introduction

The common carp (*Cyprinus carpio*) is one of the most important freshwater fish species in the world. It is a highly adaptable species that can survive in a wide range of environmental conditions. The common carp is also a popular sport fish and is often caught using fishing nets.

Managing the LAT Software

This chapter describes how the LAT software works and the tasks you must perform to implement and manage the LAT software on your system.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Starting up the LAT protocol	Section 24.4
Customizing LAT characteristics	Section 24.5
Creating a service	Section 24.5.1
Setting up ports	Section 24.5.2
Setting up printers	Section 24.5.2.1
Setting up special application services	Section 24.5.2.2
‡Enabling queued incoming requests	Section 24.5.3
Enabling outgoing LAT connections	Section 24.5.4
Managing the LATACP database size	Section 24.6
‡Alpha specific	

This chapter explains the following concepts:

Concept	Section
The LAT protocol	Section 24.1
The LAT network	Section 24.2
The LAT Control Program utility	Section 24.3

24.1 Understanding the LAT Protocol

The operating system uses the LAT (Local Area Transport) software to communicate with **terminal servers** and other systems within a local area network (LAN). Terminal servers are communication devices dedicated for connecting terminals, modems, or printers to a LAN. They offer the following features:

- Provide a cost-effective method of connecting many user terminals to a computer
- Save on cable requirements
- Maximize the number of devices that can access a computer

Managing the LAT Software

24.1 Understanding the LAT Protocol

With the LAT software, which implements the LAT **protocol**, the operating system can offer resources, or **services**, that the terminal servers can access. A system that offers LAT services is called a **service node**. In addition, nodes can access LAT services by enabling outgoing connections (using LATCP) and using the DCL command SET HOST/LAT. (In the remainder of this chapter, "servers" refers both to dedicated terminal servers and to nodes that allow outgoing access to other LAT services.)

A LAT service can consist of all the resources of a computer system, or it can be a specific resource on a computer system, such as an application program. You can set up your system as a **general timesharing service**, meaning that all of its resources are available to users in the LAN, or you can restrict access to a specific service (application program) on the system. This chapter and the *OpenVMS I/O User's Reference Manual* outline the procedure you use to set up access to a dedicated application program.

24.1.1 How the LAT Protocol Works

The LAT protocol allows the terminal servers and computers to communicate within a LAN, such as the Ethernet or the Fiber Distributed Data Interconnect (FDDI). The LAT protocol matches terminals and other devices to the computing resources (services) of the LAN. Because LAT terminals are not connected directly to the computer (service node) they are accessing, the local server must listen for service requests from its terminals and be able to match the terminals with computers that provide the desired services.

Using the LAT protocol, then, the operating system announces its available services over the LAN. Servers listen to the LAN announcements and build a database of service information so that they can locate an appropriate system when a user terminal requests computing services. For example, a user terminal might request general processing service or a data entry program on the operating system. A server uses the LAT protocol to establish and maintain a connection between the requesting terminal and the operating system.

Sometimes the operating system can request services from a terminal server. The LAT protocol allows systems to ask for connections to printers or other devices attached to a terminal server.

24.1.2 Advantages of the LAT Protocol

Using the LAT protocol on your system has many advantages:

- The LAT protocol lets you make the resources of any computer on a local area network available to any user in that network.
- In addition to general processing resources, you can set up terminals, printers, and modems so they are available from multiple systems in the local area network. This lets you efficiently use these resources and keep them available even if one of the systems in the network must be shut down.
- You can also set up application programs, such as data entry programs or news services, as resources. When a user requests a connection to the resource, the LAT protocol sets up a connection directly to the application program. No login procedure is necessary.
- The LAT protocol provides load balancing features and recovery mechanisms so users get the best, most consistent service possible. In their broadcast messages, systems rate the availability of their services so that servers can establish connections to computing resources on the least busy node. If a node

becomes unavailable for any reason, the servers attempt to provide access to alternate services.

- Users can establish multiple computing sessions on their terminals, connecting to several different computers and switching easily from one computing session to another. After switching from one session to another, users can return to the previous session and pick up where they left off. This saves users the time normally required to close out and reopen files or accounts and to return to the same point in a session.
- Finally, the LAT protocol can provide improved system performance. Because the servers bundle messages onto a single LAN interface, a server interface decreases the network traffic and reduces the number of computer interrupts realized in systems where terminals, modems, and printers each have a physical connection to the computer.

24.2 Understanding the LAT Network

A **LAT network** is any local area network where terminal servers and operating systems use the LAT protocol. A LAT network can coexist on the same LAN with other protocols. The LAT protocol, which operates on both terminal servers and the operating systems, is designed to ensure the safe transmission of data over the LAN.

The LAT network consists of the following components:

Component	For More Information
Service nodes	Section 24.2.1
Terminal server nodes	Section 24.2.2
Nodes allowing outgoing connections	Section 24.2.3
LAN cable	Section 24.2.4

Service nodes supply computing resources for the local network, while terminal server nodes (or nodes allowing outgoing connections) port their terminals, modems, or printers to those resources upon request from a user terminal or an application program.

You can use the LAT Control Program (LATCP) to configure the LAT characteristics for your system. LATCP allows you to set up your system to support:

- Incoming access only
- Outgoing access only
- Both incoming and outgoing access

The systems that support incoming LAT connections are **service nodes**. (Using LATCP, you can also set up your system so that it supports neither incoming nor outgoing access.)

Managing the LAT Software

24.2 Understanding the LAT Network

24.2.1 Service Nodes

A service node is one type of node in a LAT network. (Nodes that are not running an OpenVMS operating system can also be used along with the OpenVMS nodes in a LAT network.) A service node is an individual computer in a LAN that offers its resources to users and devices. Because the OpenVMS operating systems contain the LAT protocol, any OpenVMS system can be configured as a service node within a LAT network.

24.2.1.1 Types of Services

Each node offers its resources as a **service**. Often, a node offers a general processing service, but it can offer limited services (Alpha only) or special application services as well. Any or all of the services can be specialized applications.

For example, your service node might offer services for the following:

- General processing
- Data entry
- Stock quotations

The general processing service would allow the use of the general computing environment. The data entry and stock services, on the other hand, would be restricted environments, with connections to the application service but to no other part of the service node.

Each service is distinguished by the name the system manager assigns to it. In a VMScluster, Digital recommends that the service name be the same as the cluster name. In an independent node, Digital recommends that the service name be the same as the node name. With special service applications, the service holds the name of the application.

24.2.1.2 Service Announcements

A service node announces its services over the LAN at regular intervals so that terminal servers (and OpenVMS systems that allow outgoing connections) know about the availability of these network resources. The service announcement provides the physical node name, the service names, a description of services, and a rating of service availability. Servers listen to the LAN announcements and record information in a database. On nodes allowing outgoing connections, this database is maintained by the LAT Ancillary Control Process (LATACP). (See Section 24.6 for more information about managing the LATACP database.)

Whenever a user terminal or application program requests a service, the server node connects to the appropriate service node.

24.2.1.3 Print Requests

In some cases, service nodes can request services from terminal servers. The most common situation is when the system wants to use a printer that is connected to a terminal server port. The system submits the print request to the terminal server print queue that is set up and initialized in the OpenVMS startup procedure. Then the LAT symbiont (the process that transfers data to or from mass storage devices) requests the LAT port driver to create and terminate connections to the remote printer.

For information about setting up queues for printers connected to LAT ports, see Section 13.6.4.

24.2.2 Terminal Server Nodes

A **terminal server node** is the second type of node in a LAT network. A terminal server node is usually located near the terminals and printers it supports. The terminals and printers are physically cabled to the terminal server; the terminal server is physically connected to the LAN cable.

24.2.2.1 Locating Service Nodes

Terminal servers build and maintain a directory of services from announcements advertised over the network. Then, when terminal servers receive requests from terminal users, they can scan their service databases and locate the computer that offers the requested service.

Terminal servers not only look for the node that provides the requested service, but they can also evaluate the service rating of that node. If a requested service is offered by more than one node, then the service rating is used to select the node that is least busy. A server establishes a logical connection between the user terminal and the service node.

24.2.2.2 Setting Up Connections

One logical connection carries all the data directed from one terminal server node to a service node. That is, the server combines data from all terminals communicating with the same node onto one connection. A terminal server establishes a logical connection with a service node only if a logical connection does not already exist.

If a connection fails for any reason, a terminal server attempts to find another node offering the same service and "rolls over" the connection so users can continue their computing sessions.

Even though terminal connections are bundled together, each terminal can be uniquely identified by its name. A terminal name consists of two parts: the first part is the name of the port on the terminal server that the terminal line plugs into; the second part is the name of the terminal server node.

24.2.2.3 Servicing Nodes

Although terminal servers are usually the requesting nodes in a LAT network, sometimes service nodes request service from terminal servers. Most commonly, a service node queues print requests to remote printers connected to terminal servers.

24.2.3 Nodes Allowing Outgoing Connections

Nodes can be set up to allow incoming connections, outgoing connections, or both. Nodes (excluding those that offer incoming connections only) such as terminal server nodes can locate service nodes and set up connections. The database of information about available nodes and services is maintained by the LAT Ancillary Control Process (LATACP). (See Section 24.6 for more information about managing the LATACP database.)

On a node that is set up to allow outgoing LAT connections, a user can connect to another node in the LAT network by entering the SET HOST/LAT command. For more information, see the SET HOST/LAT command in the *OpenVMS DCL Dictionary*.

Managing the LAT Software

24.2 Understanding the LAT Network

24.2.4 Sample LAT Configuration

Figure 24-1 demonstrates the components of a LAT network. The network consists of an Ethernet cable connecting service nodes and terminal server nodes.

The three service nodes in Figure 24-1, named MOE, LARRY, and ALEXIS, each offer services to terminal server nodes on the network.

Two of the service nodes, MOE and LARRY, belong to the OFFICE cluster. (The cluster is distinguished by its computer interconnect [CI] and star coupler.) Because MOE and LARRY are clustered, their service names are the same as their cluster name. Because both service nodes offer an OFFICE service, terminal server nodes can assess the work load on both OFFICE nodes and establish a connection to a node that offers the service that is less busy.

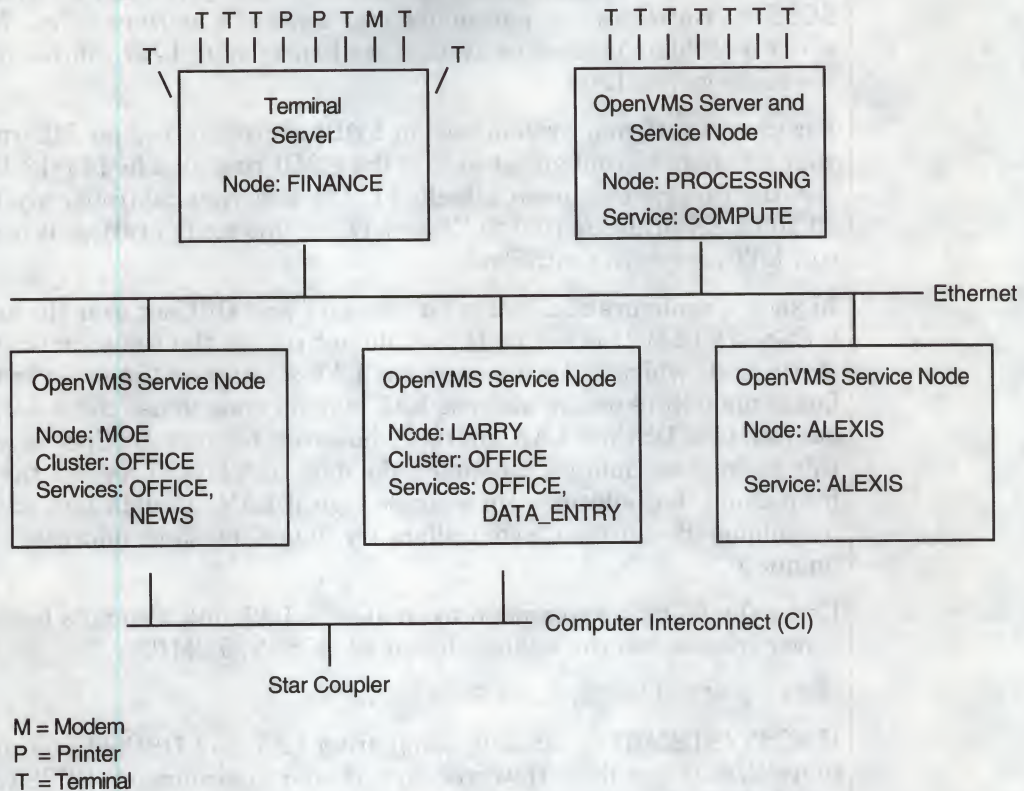
The third service node, ALEXIS, is an independent node in the LAT network so its service name is the same as its node name.

In addition to its primary OFFICE service, node MOE offers an application service called NEWS. With this specialized service, user terminals can connect directly to the online news program, without any login procedure but also without general access to the general computer resources of the node.

The node FINANCE, shown in Figure 24-1, is a terminal server node; it supports a number of interactive terminals, a modem, and a printer. The node PROCESSING is a node allowing outgoing connections; it supports several interactive terminals. The node FINANCE can accept print requests from any of the three service nodes, provided each of the service nodes has set up print queues to support remote printers on the terminal server.

Node PROCESSING is also a service node. It offers the service COMPUTE.

Figure 24-1 Sample LAT Network Configuration



24.2.5 LAT Relationship to VMScclusters and DECnet

Although the LAT protocol works independently of VMSccluster software, Digital recommends that you configure a service node to complement the VMSccluster concept. You achieve this by creating a service on each node in a VMSccluster and assigning the cluster name to this service. A terminal server assesses the availability of cluster services and establishes a connection to the node that is least busy. Thus, the LAT protocol helps balance the cluster load. If one node in the VMSccluster fails, the terminal server can transfer the failed connections to another service node within the VMSccluster.

The LAT software does not use DECnet as a message transport facility, but instead uses its own virtual circuit layer to implement a transport mechanism. The LAT and DECnet software work independently in a common LAN environment. For compatibility, if a service node is also a DECnet node, the service node name should be the same as the DECnet node name.

24.2.5.1 LAT and DECnet Running on the Same Controller

If Ethernet ports will be running both DECnet and LAT, you must start the DECnet software *before* the LAT software. If you do not start DECnet software first, all existing LAT connections may terminate, and reconnecting to the system via LAT may not be possible.

Managing the LAT Software

24.2 Understanding the LAT Network

24.2.5.2 LAT and DECnet Running on Different Controllers

If DECnet is configured on the system (or if the system is part of a cluster), the SCSSYSTEMID system parameter may contain a nonzero value. Normally, this is not a problem unless the system has two or more LAN controllers connected to the same logical LAN.

For example, if your system has an FDDI controller and an Ethernet controller, your site may be configured so that the FDDI ring attached to the FDDI controller and the Ethernet segment attached to the Ethernet controller are bridged by a 10/100 LAN bridge (FDDI-to-Ethernet). In this configuration, it is impossible to run LAT over both controllers.

In such a configuration, you *must* run LAT and DECnet over the same controller if SCSSYSTEMID is not 0. If they do not run on the same controller, DECnet starts first, which in turn causes the LAT startup on the other controller to fail. This failure occurs because LAT startup tries to use the AA-00-04-00-xx-xx address (the DECnet LAN address); however, because DECnet is already using this address on another controller, the data link layer prevents the LAT startup from using that address. (In a single logical LAN, all data link addresses must be unique. Because both controllers try to use the same address, it is no longer unique.)

Using the following command to create the LAT link also fails because the LAN driver tries to use the address based on SCSSYSTEMID:

```
LATCP> CREATE LINK LAT$LINK_2 /NODECNET
```

If SCSSYSTEMID is set to 0, configuring LAT and DECnet on different controllers is possible. However, in a cluster environment, SCSSYSTEMID cannot be set to 0.

24.3 Understanding the LATCP Utility

The LAT Control Program (LATCP) utility is a utility program used for configuring and controlling the LAT software on OpenVMS systems. LATCP commands let you stop and start the LAT driver (which implements the LAT protocol) and modify or display LAT characteristics of the OpenVMS node.

With LATCP, you can set up your system as a service node, which offers one or more resources (services) for access by users on other systems in the local area network (LAN).

In addition to being able to set up your system to allow users on other systems to access its services, you can also use LATCP to set up the system to allow its users to access services on other systems in the LAN. In this case, the system can act like a terminal server: it can manage multiple user sessions simultaneously for connections to services on other nodes.

You can use LATCP to set up your system to support incoming access only, outgoing access only, or both incoming and outgoing access. You can also set up your system so that it supports neither incoming nor outgoing access.

When you set up your system to support outgoing access, the LAT software manages a database of LAT services and nodes. The software builds the database when you enable outgoing access on your node. The software begins to collect **LAT service announcements**—multicast messages sent by LAT service nodes—and builds the database based on these service announcements. You can use LATCP to display the services and nodes in this database and to control the size

of the database. Allow outgoing access on systems that can tolerate the additional overhead, such as standalone systems.

Use LATCP to do the following:

- Specify operational characteristics for your node and its services
- Turn the state of the LAT port driver (LTDRIVER) on and off
- Display the status of LAT services and service nodes in the network
- Display the status of links created on your LAT node
- Display the status of your LAT node
- Show and zero LAT counters
- Create, delete, and manage LAT ports
- Recall previously entered LATCP commands so that you can execute them again without having to reenter them
- Create subprocesses so that you can execute DCL commands without exiting from LATCP

With the LAT protocol, you can set up LAT application ports on the local node so that users can access printers and other asynchronous devices that are connected to LAT terminal servers or service nodes on the LAN. The remote devices must be configured appropriately.

24.3.1 Invoking and Exiting LATCP

Enter the following command to invoke LATCP:

```
$ RUN SYS$SYSTEM:LATCP
LATCP>
```

At the LATCP> prompt, you can enter LATCP commands. To exit LATCP, enter EXIT or press Ctrl/Z at the LATCP> prompt.

You can also execute a single LATCP command by using a DCL string assignment statement, as shown in the following example:

```
$ LCP ::= $LATCP
$ LCP SET NODE/STATE=ON
```

LATCP executes the SET NODE command and returns control to DCL.

24.3.2 LATCP Commands

Table 24–1 summarizes the LATCP commands.

Table 24–1 LATCP Commands

Command	Function
ATTACH	Transfers control from your current process to the specified process.
CREATE LINK	Creates LAT data links.
CREATE PORT	Creates a logical port on the local node.
CREATE SERVICE	Creates a service on a service node.

(continued on next page)

Managing the LAT Software

24.3 Understanding the LATCP Utility

Table 24–1 (Cont.) LATCP Commands

Command	Function
DEFINE/KEY	Assigns a command string to a function key on your keypad.
DELETE LINK	Deletes a LAT data link from a node.
DELETE PORT	Deletes an application port or dedicated port.
‡ DELETE QUEUE_ENTRY	Deletes an incoming queued request from the local node.
DELETE SERVICE	Deletes a service on a service node.
EXIT	Returns the user to DCL command level.
HELP	Displays help text for LATCP commands.
RECALL	Recalls LATCP commands that you entered previously so that you can execute them again.
REFRESH	Refreshes your display screen, for example, after your display has been overwritten by output from some other source.
‡ SCROLL	Allows you to retrieve information that has scrolled off the screen.
SET LINK	Modifies characteristics of LAT data links.
SET NODE	Specifies LAT characteristics for a node.
SET PORT	Maps a logical port on a node to either a remote device on a terminal server or a special application service on a remote LAT service node.
SET SERVICE	Changes service characteristics.
SHOW LINK	Displays the characteristics of links on your node.
SHOW NODE	Displays the characteristics of nodes.
SHOW PORT	Displays port characteristics.
‡ SHOW QUEUE_ENTRY	Displays information about requests, or entries, queued on the local node.
SHOW SERVICE	Displays characteristics of LAT services known to your node.
SPAWN	Creates a subprocess.
ZERO COUNTERS	Resets the node counters, service counters, and link counters maintained by your node.

‡Alpha specific

For detailed information about LATCP commands and qualifiers, see the *OpenVMS System Management Utilities Reference Manual*.

24.4 Starting Up the LAT Protocol

As system manager, you start up the LAT protocol and configure your node as a service node by executing the command procedure `SYS$STARTUP:LAT$STARTUP`. This procedure executes the following two procedures:

1. `LAT$CONFIG.COM`, to load the LAT terminal driver `LTDRIVER` and create the `LATACP` process

Managing the LAT Software

24.4 Starting Up the LAT Protocol

2. LAT\$SYSTARTUP.COM, to execute LATCP commands that define LAT characteristics

How to Perform This Task

To make sure the LAT protocol is started each time the system boots, add a command to execute this procedure in the general-purpose, site-specific startup command procedure, described as follows. (See Section 5.2.1 for more detailed information about this command procedure, including the file specification used to identify it in your operating system.)

To set up your node as a LAT service node and start the LAT protocol software on your system each time the system boots, edit the general-purpose, site-specific startup command procedure to add the following line:

```
$ @SYS$STARTUP:LAT$STARTUP.COM
```

When the general-purpose, site-specific startup command procedure executes this command, it invokes LAT\$STARTUP.COM, which in turn invokes the LAT\$CONFIG and LAT\$SYSTARTUP command procedures.

You can append any of the following arguments to the command line that invokes LAT\$STARTUP to specify unique LAT characteristics for your node. The procedure will pass these arguments to LAT\$SYSTARTUP.COM to define the LAT characteristics you specify.

```
$ @SYS$STARTUP:LAT$STARTUP "P1" "P2" "P3" "P4" "P5"
```

Digital recommends that you modify LAT\$SYSTARTUP.COM directly, rather than passing parameters in *P1* through *P5*. However, if you choose to use *P1* through *P5*, the arguments have the following meanings:

Argument	Format	Meaning
<i>P1</i>	Service-name	Name of the service. For clustered service nodes, use the cluster alias as the service name. For independent service nodes, use the DECnet node name. LAT\$SYSTARTUP.COM uses the argument <i>P1</i> to assign a service name to the node (with the LATCP CREATE SERVICE command).
<i>P2-P4</i>	Any of the following: /IDENTIFICATION= "string" /GROUPS=(ENABLE=group-list)	LAT\$SYSTARTUP.COM uses the arguments to assign LAT node characteristics (with the LATCP SET NODE command). Description of the node and its services that is advertised over the local area network (LAN). The default is the string defined by the logical name SYS\$ANNOUNCE. Make sure you include five sets of quotation marks around the identification string. For example: "/IDENTIFICATION=" - """"Official system center"""" Terminal server groups qualified to establish connections with the service node. By default, group 0 is enabled.

Managing the LAT Software

24.4 Starting Up the LAT Protocol

Argument	Format	Meaning
	/GROUPS=(DISABLE= <i>group-list</i>)	Removes previously enabled terminal server groups. If you are specifying the preceding qualifier to enable groups, you can combine the qualifiers into one, as shown in the example that follows this table.
P5	Any qualifiers valid with the CREATE SERVICE command	LAT\$SYSTARTUP.COM uses this argument to assign service characteristics with the LATCP CREATE SERVICE command. You can specify the /IDENTIFICATION, /LOG, and /STATIC_RATING qualifiers. Specify several qualifiers as shown in the following example: <pre>"/IDENTIFICATION=" - "****Official system node****" - /STATIC_RATING=250"</pre>

Note that if you want to do any of the following LAT network tasks, you must edit LAT\$SYSTARTUP.COM (described in Section 24.5):

- Set up LAT printers.
- Create special application services.
- Set up the node to allow outgoing connections (to support the SET HOST/LAT command).

For a full description of LATCP commands and qualifiers, see the *OpenVMS System Management Utilities Reference Manual*.

Example

The following command creates the service OFFICE on the service node MOE, which is part of the OFFICE cluster (refer to Figure 24-1):

```
$ @SYS$STARTUP:LAT$STARTUP OFFICE
```

24.5 Customizing LAT Characteristics

To define special LAT characteristics for your node, edit the site-specific command procedure SYS\$MANAGER:LAT\$SYSTARTUP.COM. This command procedure contains LATCP commands that define LAT characteristics. LAT\$SYSTARTUP.COM is invoked when you execute the LAT\$STARTUP command procedure. As explained in Section 24.4, you typically execute LAT\$STARTUP.COM from the general-purpose, site-specific startup command procedure.

If you want your node to be a LAT service node that only supports incoming connections from interactive terminals, editing LAT\$SYSTARTUP.COM is not necessary. You can assign a service name and other characteristics by specifying parameters when you invoke the command procedure SYS\$STARTUP:LAT\$STARTUP, as described in Section 24.4.

However, you can edit LAT\$SYSTARTUP.COM to add LATCP commands that customize LAT characteristics for your node, for example:

Managing the LAT Software

24.5 Customizing LAT Characteristics

Task	For More Information
Create more than one service	Section 24.5.1
Create logical ports for special application services and printers	Section 24.5.2
‡Enable queued incoming requests	Section 24.5.3
Enable outgoing LAT connections to support the SET HOST/LAT command	Section 24.5.4
Tailor node characteristics ¹	Section 24.5.5

¹For example, to assign special service announcements or LAN links (using the SET NODE and SET LINK commands).

‡Alpha specific

Caution

Do not edit the command procedures LAT\$STARTUP.COM and LAT\$CONFIG.COM. These are procedures supplied by Digital to perform functions necessary for the LAT protocol to run correctly. Edit only LAT\$SYSTARTUP.COM to define LAT characteristics specific to your site.

If you edit LAT\$SYSTARTUP.COM, you should add only LATCP commands. In addition, you should conform to the order of commands in the template file SYS\$MANAGER:LAT\$SYSTARTUP.TEMPLATE. Section 24.5.5 provides a sample edited LAT\$SYSTARTUP procedure. The *OpenVMS System Management Utilities Reference Manual* contains full descriptions of all the LATCP commands you can include in LAT\$SYSTARTUP.COM.

24.5.1 Creating Additional Services

The LAT\$SYSTARTUP.COM procedure provided by Digital creates one service. This can be a primary service, one through which users can access the general computing environment. It can also be a special application service, such as a data entry program or an online news service.

Alpha

On Alpha systems, you can also create a limited service with a fixed number of LTA devices, as described in Section 24.5.2.3. ♦

The LAT\$SYSTARTUP.COM procedure creates the service with the same name as that of your node, unless you specify a unique service name as an argument to the @SYS\$STARTUP:LAT\$STARTUP.COM command, as explained in Section 24.4.

How to Perform This Task

To create services in addition to the one provided in LAT\$SYSTARTUP.COM, use the CREATE SERVICE commands, which you can add to LAT\$SYSTARTUP.COM. Note that if you create an application service, Digital recommends that you assign the name of the application program. For more information on the LATCP command CREATE SERVICE, see the *OpenVMS System Management Utilities Reference Manual*.

Managing the LAT Software

24.5 Customizing LAT Characteristics

Example

The following example creates the application service NEWS on the local node.

```
$ LCP ::= $LATCP
$ LCP CREATE SERVICE /APPLICATION NEWS
```

24.5.2 Setting Up Ports

The LAT\$SYSTARTUP.COM file provided by Digital includes sample commands to create logical ports on the service node and associates them with physical ports or services on the terminal server node. These ports can be used for application services and remote printers.

How to Perform This Task

To create ports, enable the sample commands by removing the exclamation points (!) that precede them in the LAT\$SYSTARTUP.COM file, or add similar CREATE PORT and SET PORT commands to that file to meet your needs. For information on the LATCP commands CREATE PORT and SET PORT, see the *OpenVMS System Management Utilities Reference Manual*.

Note

Digital strongly recommends that you create application and dedicated ports *after* the LATCP command SET NODE/STATE=ON is executed. This minimizes nonpaged pool memory usage and eliminates the possibility of creating duplicate ports.

Note that you may encounter the following error when attempting to create an application port (with a command such as LCP CREATE PORT LTA5001: /APPLICATION, for example):

```
%LAT-W-CMDERROR, error reported by command executor
-SYSTEM-F-DUPLNAM, duplicate name
```

This error indicates that the LAT application port you are trying to create is already created by some other application. This application could be LATCP itself (LATCP's port, LATCP\$MGMT_PORT, is used to communicate with LTDRIVER).

To avoid this error, make sure the SET NODE/STATE=ON command is executed before any commands that create application or dedicated ports. You can also use the LATCP command SET NODE/DEVICE_SEED. For more information on the SET NODE/DEVICE_SEED command, see the *OpenVMS System Management Utilities Reference Manual*.

24.5.2.1 Setting Up Printers

If you set up a port for a printer, you must also perform the following tasks:

1. Create a spooled output queue for the printer.
2. Add a command to start the queue to the startup command procedure that starts your queues, or to the general-purpose, site-specific startup command procedure.

These tasks are described in Chapter 13.

Managing the LAT Software

24.5 Customizing LAT Characteristics

24.5.2.2 Setting Up Special Application Services

To establish a special application service, include the /DEDICATED qualifier when defining a LAT port. The application program to which the service connects must define the same dedicated port. For example, the following commands set up ports for an application service called NEWS:

```
$ LCP ::= $LATCP
$ LCP CREATE PORT LTA333:/DEDICATED
$ LCP SET PORT LTA333:/SERVICE=NEWS
```

Before application services can be available to user terminals on the LAT network, you must start the application program. You usually add commands to SYLOGIN.COM to do this.

24.5.2.3 Setting Up Limited Services

Application services with dedicated ports allow you to create a predetermined number of LTA devices (LAT terminals, for example) that are under the control of a process supplied by the system. In that environment, however, the user cannot log in to the service because no way exists for dedicated LTA devices to run the system login image (LOGINOUT.EXE).

You can create a **limited** service that allows users to log in to a predetermined number of LTA devices associated with that limited service. When all those devices are in use, the LAT software will reject additional connection requests to that service, as indicated by "service in use" error messages. Creating a limited service in this way allows you to control the number of LAT users on your system. (Note, however, that you cannot control which LTA device will be assigned when a user connects to the limited service.)

The following example sets up a limited service with two predetermined LTA devices:

```
$ LCP ::= $LATCP
$ LCP CREATE SERVICE /LIMITED RESTRICTED
$ LCP CREATE PORT LTA100 /LIMITED
$ LCP CREATE PORT LTA101 /LIMITED
$ LCP SET PORT LTA100 /SERVICE=RESTRICTED
$ LCP SET PORT LTA101 /SERVICE=RESTRICTED
```

When a user attempts to connect to the limited service named RESTRICTED, the LAT software will choose either LTA100 or LTA101 (whichever is available first) and complete the user connection. The user can then log in to that system. If another user connects to the service, that second connection attempt will be assigned to the remaining LTA device. The user can then log in to that second system. When the two devices associated with the limited service named RESTRICTED are both in use, any subsequent attempts to connect to that service will be rejected, as indicated by the "service in use" error message.

When a user logs out of the system (LTA100 or LTA101), that LTA device is *not* deleted. Instead, it is reset to accept the next connection request to the limited service.

24.5.3 Queuing Incoming Requests

By default, incoming requests to limited or application services are queued. This means that if you attempt to connect to a limited or application service (by using a terminal server port with forward queuing enabled or by entering the DCL command SET HOST/LAT/QUEUE), the LAT software will queue, rather than reject, this connection request if the service has no available ports.

Managing the LAT Software

24.5 Customizing LAT Characteristics

How to Perform This Task

You can set up and manage a service that queues incoming connect requests as follows:

- Use the LATCP command SHOW SERVICE to determine whether the service has queuing enabled or disabled.
- If queuing is disabled, use the SET SERVICE /QUEUE command to enable queuing.
- Use the SET NODE /QUEUE_LIMIT=*n* command on the local node to control the number of free queue slots (where *n* is between 0 and 200).
- Use the SHOW QUEUE_ENTRY [entry-id] command to view entries in the local queue.
- Use the DELETE QUEUE_ENTRY [entry-id] command to delete an entry from the local queue.

See the *OpenVMS System Management Utilities Reference Manual* for more detailed descriptions of the LATCP commands and qualifiers you use to support queued requests.

Example

The following example shows how to enable queuing on your system:

```
$ LCP :== $LATCP
$ LCP SET SERVICE /QUEUE
```

Note

If a system is configured to handle queued connect requests, that system *must* be set up as follows to avoid possible queue connection failures:

- Incoming and outgoing connections must be enabled.
 - User group codes and service group codes must be identical.
-

24.5.4 Enabling Outgoing LAT Connections

By default, outgoing LAT connections are disabled on a node. If you want to allow users to use the SET HOST/LAT connection to establish LAT connections from the node, you must edit LAT\$SYSTARTUP.COM to enable outgoing connections. For more details on using the SET HOST/LAT command for outgoing LAT connections, see the description of that command in the *OpenVMS DCL Dictionary*.

Commands to enable outgoing connections are included in the LAT\$SYSTARTUP.COM file provided by Digital. Enable the command of your choice by removing the exclamation point (!) that precedes it, or add a similar command to meet your needs. For more information, see the /CONNECTIONS and /USER_GROUPS qualifiers to the LATCP command SET NODE in the *OpenVMS System Management Utilities Reference Manual*.

To attain optimal SET HOST/LAT performance and forward port performance, you should set the system parameter TTY_ALTYPAMD to 1500 and reboot.

Managing the LAT Software

24.5 Customizing LAT Characteristics

If you want to set up your node only as a service node with incoming connections enabled, editing LAT\$SYSTARTUP.COM is not necessary. However, you might edit LAT\$SYSTARTUP.COM to do one or more of the following tasks:

- Create more than one service on a node
- Create special application services
- Set up LAT printers
- Enable outgoing LAT connections (to allow a node to act as a terminal server node)
- Tailor node characteristics; for example, to assign special service announcements or connections to the LAN

24.5.5 Sample Edited LAT\$SYSTARTUP.COM Procedure

The following is a sample of an edited LAT\$SYSTARTUP.COM procedure that creates services, creates and sets ports, and sets nodes to allow incoming and outgoing connections.

```
$! Copyright (c) 1992 Digital Equipment Corporation. All rights reserved.
$!
$! LAT$SYSTARTUP.COM -- LAT Startup Commands Specific to Site
$!
$! Use this command procedure to customize the LAT characteristics for
$! the local node. These commands, which should serve as examples,
$! will set up a LAT service name SYS$NODE and default identification
$! SYS$ANNOUNCE. The LAT service name and identification will default
$! to SYS$NODE and SYS$ANNOUNCE unless you specify a service name and
$! identification as arguments to the command line that invokes
$! LAT$STARTUP.COM:
$!      $ @SYS$STARTUP:LAT$STARTUP
$!
$! You can specify other node and service characteristics (such as group
$! codes) as arguments to this command line, as shown below.
$!
$!      Argument      Function
$!      -----      -
$!
$!      P1             Name of the service to be created. If not supplied, a
$!                     service will be created with the same name as the node.
$!
$!      P2,P3,P4       Parameters and qualifiers to the SET NODE command.
$!
$!      P5             Parameters and qualifiers to the SET SERVICE command.
$!                     P5 is only used if P1 is specified. More than one
$!                     argument may be supplied by enclosing the string in
$!                     quotes.
$!
$!      Example: $ @SYS$STARTUP:LAT$STARTUP HAWK "/IDENTIFICATION=" -
$!      """"Development node""""
$!
$!      Please review and edit this file for possible additions and deletions
$!      that you wish to make. Future software updates will not overwrite the
$!      changes made to this file.
$!
$!      required_privileges = "OPER"
$!      prev_privs = f$setprv(required_privileges)
$!      if .not. f$privilege(required_privileges) then goto no_privileges
$!      lcp := $latcp
$!
$! ----- Modify Node Characteristics -----
$!
```


Managing the LAT Software

24.5 Customizing LAT Characteristics

```
$ lcp set node 'p2' 'p3' 'p4'
$!
$! Some examples:
$!
$! ** Allow incoming connections only
$!
$! lcp set node /connections=incoming /groups=(enable=(12,40,43,73),disable=0)
$! lcp set node /connections=incoming /groups=enable=(0-255)
$!
$ LCP SET NODE /CONNECTIONS=INCOMING /GROUPS=(ENABLE=(12,40,43,73),DISABLE=0)
$!
$! ** Allow outgoing connections only
$!
$! lcp set node /connections=outgoing /user_groups=enable=(24,121-127)
$! lcp set node /connections=outgoing /user_groups=(enable=0-255) /node_limit=50
$!
$! ** Enable incoming and outgoing connections
$!
$! lcp set node /connections=both /group=enable=(43,73) /user=enable=(44,56)
$! lcp set node /connections=both /group=enable=(0-255) /user=enable=(0-255)
$!
$!
$! ----- Modify Service Characteristics -----
$!
$ if p1 .eqs. ""
$ then
$   lcp create service
$ else
$   lcp create service 'p1' 'p5'
$ endif
$! ----- Start LAT Protocol -----
$!
$ lcp set node /state=on
$!
$!
$! ----- Create and Map Ports -----
$!
$! Some examples:
$!
$! lcp create port lta101: /dedicated
$! lcp create port lta102: /application
$! lcp create port lta103: /application
$! lcp create port /nolog/logical=(name=ln03$mgmt, table=system, mode=executive)
$
$ LCP CREATE PORT LTA1: /NOLOG
$ LCP CREATE PORT LTA20: /NOLOG
$
$! lcp set port lta101: /dedicated /service=graphics
$! lcp set port lta102: /node=server_1 /port=port_1
$! lcp set port lta103: /node=server_2 /service=laser
$! lcp set port ln03$mgmt: /node=server_3 /service=ln03_printers
$!
$ LCP SET PORT LTA1: /APPLICATION/NODE=TERM_SERVER_1 /PORT=PORT_6
$ LCP SET PORT LTA20: /APPLICATION/NODE=TERM_SERVER_2 /PORT=PORT_6
$!
$exit:
$ prev_privs = f$setprv(prev_privs)
$ exit
$!
$no_privileges:
$ write sys$output "Insufficient privileges to execute LATCP commands."
$ write sys$output "Requires ",required_privileges," privileges."
$ goto exit
```


24.6 Managing the LATACP Database Size

On OpenVMS nodes, another component of the LAT software, the LAT Ancillary Control Process (LATACP), maintains the database of available nodes and services. The nodes and services can be those that are multicast from remote LAT nodes, or they can consist of the local node and one or more local services that you create on your own system. The maximum size of this database is dependent on the value of the system parameter CTLPAGES.

After you enter a LATCP command, you might get the following response:

```
%LAT-W-CMDERROR, error reported by command executor  
-LAT-F-ACPNOCTL, insufficient resources - ACP CTL/P1 space limit reached
```

If so, this signifies that the database size has reached the CTLPAGES limit. You can correct the situation in one of three ways:

- Reduce the size of the database by reducing the node limit. Use the LATCP command SHOW NODE to display the node limit; use the command SET NODE/NODE_LIMIT to change it. For more information, see the *OpenVMS System Management Utilities Reference Manual*.
- Reduce the size of the database by reducing the user group codes that are enabled on the node. Use the LATCP command SHOW NODE to display the enabled user group codes; use the command SET NODE/USER_GROUPS=DISABLE to disable some of them. For more information, see the *OpenVMS System Management Utilities Reference Manual*.

If you choose this step, you must also edit your startup procedures to change the user groups that are enabled each time the system reboots. For more information, see Section 24.5.

- Extend the size of the database by increasing the value of the system parameter CTLPAGES. As a general rule, note that every unit of CTLPAGES that you increase is roughly equivalent to six additional nodes or services that will be stored in the database.

After you change CTLPAGES, you must reboot the system for the changed value to take effect. Make sure you add the increased value of CTLPAGES to the AUTOGEN parameter file MODPARAMS.DAT. For more information on changing values of system parameters, see Section 14.2.

578-11.1000 THE LABOR DIVISION OF THE BUREAU OF ECONOMIC ANALYSIS

The Labor Division of the Bureau of Economic Analysis (BEA) is responsible for the collection, analysis, and dissemination of data on the labor market. The division's primary focus is on the nonfarm business sector, which accounts for approximately 80 percent of the U.S. economy. The division's data are used by a wide range of users, including government agencies, academic researchers, and the private sector.

The division's data are collected from a variety of sources, including the Current Population Survey (CPS), the Survey of Current Business (SCB), and the Survey of Professional and Technical Personnel (SPTP).

The division's data are analyzed using a variety of statistical techniques, including time series analysis, cross-sectional analysis, and regression analysis.

The division's data are disseminated through a variety of channels, including the BEA's website, the BEA's publications, and the BEA's data services.

The division's data are used to monitor the labor market and to provide information on the economic conditions of the United States.

The division's data are also used to provide information on the economic conditions of the United States to the public.

The division's data are also used to provide information on the economic conditions of the United States to the private sector.

The division's data are also used to provide information on the economic conditions of the United States to academic researchers.

The division's data are also used to provide information on the economic conditions of the United States to government agencies.

The division's data are also used to provide information on the economic conditions of the United States to the public.

The division's data are also used to provide information on the economic conditions of the United States to the private sector.

The division's data are also used to provide information on the economic conditions of the United States to academic researchers.

The division's data are also used to provide information on the economic conditions of the United States to government agencies.

The division's data are also used to provide information on the economic conditions of the United States to the public.

The division's data are also used to provide information on the economic conditions of the United States to the private sector.

The division's data are also used to provide information on the economic conditions of the United States to academic researchers.

The division's data are also used to provide information on the economic conditions of the United States to government agencies.

The division's data are also used to provide information on the economic conditions of the United States to the public.

Managing DECdtm Services

This chapter describes what you must do if you want to run software that uses DECdtm services. Software products that can currently use DECdtm services include ACMS, DBMS, DECintact, Rdb, and RMS Journaling.

Information Provided in This Chapter

This chapter describes the following tasks:

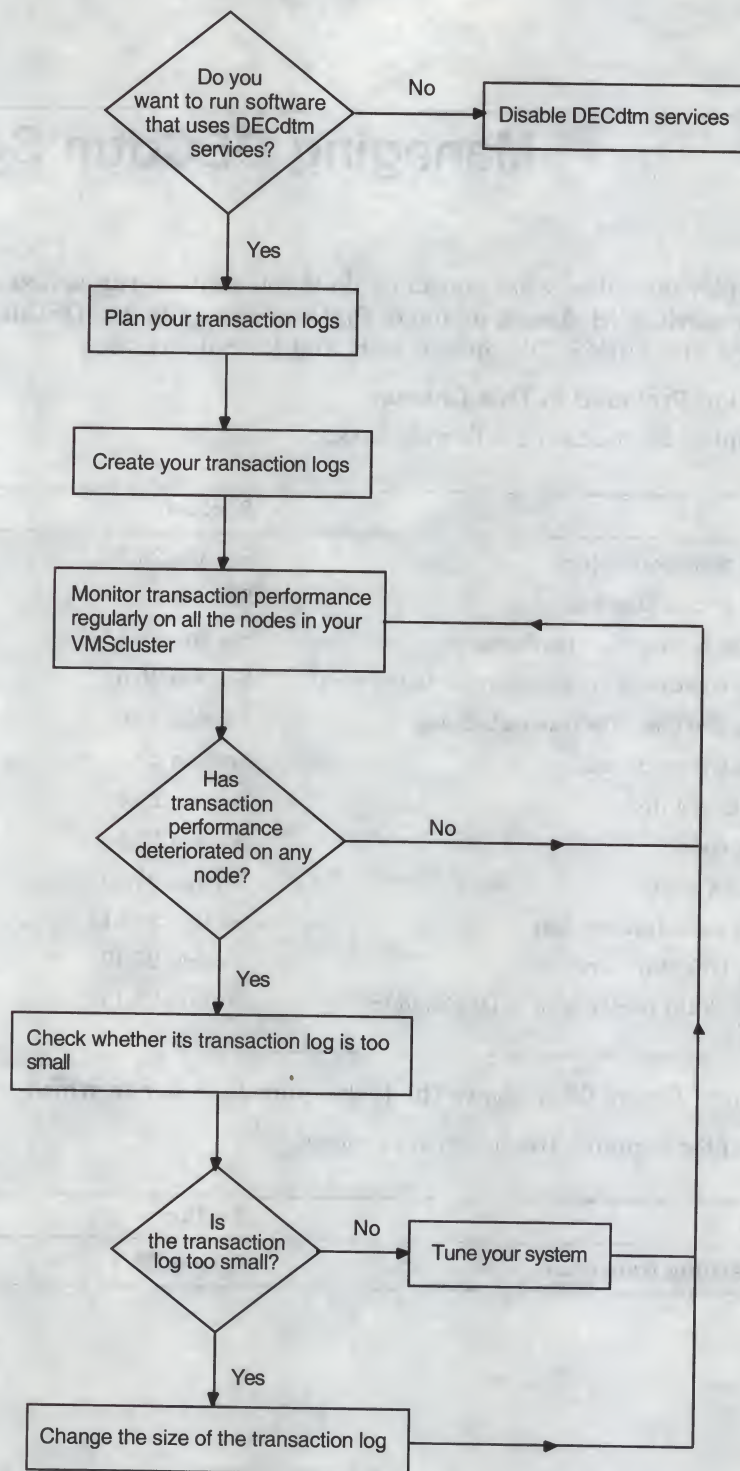
Task	Section
Planning transaction logs	Section 25.2
Creating transaction logs	Section 25.3
Monitoring transaction performance	Section 25.4
Checking whether a transaction log is too small	Section 25.5
Changing the size of a transaction log	Section 25.6
Moving a transaction log	Section 25.7
Dismounting a disk	Section 25.8
Adding a node	Section 25.9
Removing a node	Section 25.10
Disabling DECdtm services	Section 25.11
Enabling DECdtm services	Section 25.12
Using DECdtm Services in a DECnet/OSI Network	Section 25.13

The map in Figure 25-1 shows the tasks, and the order in which to do them.

This chapter explains the following concept:

Concept	Section
Understanding transaction logs	Section 25.1

Figure 25-1 Managing DECdtm Services



ZK-5154A-GE

25.1 Understanding Transaction Logs

A **transaction log** is a file that stores information about DECdtm transactions performed on a node. It is of file type LM\$JOURNAL.

Before a node can execute DECdtm transactions, you must create a transaction log for the node. In a VMScLuster, create a transaction log for each node. Use the Log Manager Control Program (LMCP) utility to create and manage transaction logs.

DECdtm services use the logical name SYS\$JOURNAL to find transaction logs. You must define SYS\$JOURNAL to point to the directories that contain transaction logs.

25.2 Planning Transaction Logs

The size and location of a transaction log can affect transaction performance. Before you create a transaction log, decide the size and location of the transaction log.

Later, you can change the size of a transaction log, or move it. However, careful planning at this stage reduces the need for future changes.

This section describes:

Task	Section
Deciding the size of a transaction log	Section 25.2.1
Deciding the location of a transaction log	Section 25.2.2

25.2.1 Deciding the Size of a Transaction Log

When you create a transaction log, you can specify its size. The default size is 4000 blocks; this gives acceptable performance on most systems.

If you know the expected rate of transactions, Digital suggests the following formula to calculate the transaction log size:

$$size = 40 * rate$$

where:

size is the size of the transaction log in blocks.

rate is the average number of transactions executed per second.

If you do not know the rate of transactions, accept the default size of 4000 blocks.

25.2.2 Deciding the Location of a Transaction Log

If possible, choose a disk that is:

Fast	Achieve speed by using a high-performance disk, such as a solid-state disk, that is not heavily used.
Highly available	Achieve high availability by having multiple access paths to the data. In a VMScluster environment, use a disk that can be accessed by the other nodes in the cluster. This ensures that if one node fails, transactions running on other nodes are not blocked.
Reliable	Achieve reliability by keeping multiple copies of the data. Using a shadowed disk is more reliable than using a nonshadowed disk, but may be slower because transaction logs are almost exclusively write-only.

You may need to choose between speed and either availability or reliability. For example, if the node is a workstation, you may choose to sacrifice speed for availability and reliability by putting the node's transaction log on a shadowed HSC-based disk, instead of on a faster disk attached to the workstation.

In a VMScluster environment, if possible distribute the transaction logs across different disks. Having more than one transaction log on a disk can lead to poor transaction performance.

Note

Make sure that the disk has enough contiguous space to hold the transaction log. A discontinuous transaction log leads to poor transaction performance.

25.3 Creating Transaction Logs

Before a node can perform DECdtm transactions, you must create a transaction log for the node. In a VMScluster environment, create a transaction log for each node.

Caution

Removing a node from a VMScluster after you have created the transaction logs can lead to data corruption. For instructions on how to remove a node safely, see Section 25.10.

How to Perform This Task

1. For each node, decide the size and location of the transaction log, using the guidelines in Section 25.2. Remember that the disks must have enough contiguous space to hold the transaction logs.
2. If you are in a VMScluster environment, make sure that the disks on which you want to create the transaction logs are mounted clusterwide.

If your VMScluster system operates in a DECnet/OSI network, you must include a node's SCSNODE name in the name of the transaction log for that node. A node's SCSNODE name is defined by the SCSNODE system parameter.

Managing DECdtm Services

25.3 Creating Transaction Logs

3. Decide in which directories you want to create the transaction logs. You may want to create new directories for the transaction logs.
4. Define SYS\$JOURNAL to point to the directories in which you want to create the transaction logs:

```
DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL dirspec[,...]
```

where *dirspec* is the full specification of a directory in which you want to create one or more transaction logs. List all the directories that will contain transaction logs. You can list the directories in any order.

In a VMScLuster environment, use SYSMAN to define SYS\$JOURNAL clusterwide.

5. Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to include the SYS\$JOURNAL definition.

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

6. Create one transaction log for each node, using LMCP's CREATE LOG command:

```
CREATE LOG [/SIZE=size] dirspecSYSTEM$node.LM$JOURNAL
```

where:

size is the size of the transaction log in blocks. By default, the size of the transaction log is 4000 blocks.

dirspec is the full specification of the directory in which you want to create the transaction log.

node is the name of the node.

7. Make sure DECdtm services are enabled as follows:

Step	Action
a.	Check whether the logical SYS\$DECDTM_INHIBIT is defined: \$ SHOW LOGICAL SYS\$DECDTM_INHIBIT
b.	Is SYS\$DECDTM_INHIBIT defined? Yes DECdtm services are disabled. Enable DECdtm services by following the instructions in Section 25.12. No DECdtm services are enabled.

Example

This example shows how to create transaction logs for nodes in a VMScLuster system and whose SCSNODE names are BLUE and RED. Neither node has a node-specific version of SYLOGICALS.COM.

Decide the size and location of the transaction logs:

Node	Size of Log (in Blocks)	Disk
BLUE	5000	DUA1
RED	4000	DUA2

Managing DECdtm Services

25.3 Creating Transaction Logs

Mount the disks clusterwide:

```
$ MOUNT/CLUSTER/SYSTEM DUA1: LOG1
$ MOUNT/CLUSTER/SYSTEM DUA2: LOG2
```

Create directories for the transaction logs:

```
$ CREATE/DIRECTORY DISK$LOG1:[LOGFILES]
$ CREATE/DIRECTORY DISK$LOG2:[LOGFILES]
```

Define SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL -
SYSMAN> DISK$LOG1:[LOGFILES], DISK$LOG2:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to include the following line:

```
$ !
$ DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL DISK$LOG1:[LOGFILES], DISK$LOG2:[LOGFILES]
$ !
```

Create the transaction logs:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CREATE LOG/SIZE=5000 DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
LMCP> CREATE LOG DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL
LMCP> EXIT
```

Make sure DECdtm services are enabled:

```
$ SHOW LOGICAL SYS$DECDTM_INHIBIT
%SHOW-S-NOTRAN, no translation for logical name SYS$DECDTM_INHIBIT
```

SYS\$DECDTM_INHIBIT is undefined, so DECdtm services are enabled.

25.4 Monitoring Transaction Performance

Changes to your system, such as increase in workload, can affect transaction performance. Once a month, monitor transactions on the node to make sure that transaction performance has not deteriorated. In a VMScluster environment, monitor transaction performance on all the nodes.

How to Perform This Task

1. Monitor transactions, using the MONITOR TRANSACTION command:

```
MONITOR TRANSACTION/SUMMARY[=summary-file]/ENDING=end-time/NODE=node[,...]
```

where:

summary-file

is the file specification of the summary file. Information about transactions is summarized and recorded in the summary file. If you omit the file specification, the information is recorded in MONITOR.SUM in your default directory.

end-time

is the time that the monitoring session ends.

node

is the name of a node. In a VMScluster environment, list all the nodes in the VMScluster.

For the best results, monitor transactions for a day at a time.

You can monitor transactions in batch mode by including the MONITOR TRANSACTION command in a command procedure.

Managing DECdtm Services

25.4 Monitoring Transaction Performance

For a full description of the MONITOR TRANSACTION command, see the *OpenVMS System Management Utilities Reference Manual*.

2. Examine the summary file.

The summary file contains values for a number of different data items. Note the following values for each node:

- Average end rate. This is the average number of transactions completed per second.
- Average completion rates. These are the average numbers of transactions completed in the following times:
 - Less than 1 second
 - Between 1 and 2 seconds
 - Between 2 and 3 seconds
 - Between 3 and 4 seconds
 - Between 4 and 5 seconds
 - More than 5 seconds

Keep a note of these values.

3. Compare the results from this monitoring session with the results from previous sessions.

For the same work load, the rate and duration of transactions should remain about the same. Indications of performance deterioration are:

- The average end rate decreases
- The average duration increases

To find out whether the average duration of transactions has increased, compare the average completion rates. If a greater proportion of the transactions takes longer to complete, the average duration of transactions has increased.

Note any trends over a number of monitoring sessions. Variations from one monitoring session to the next are probably due to variations in work load.

If you suspect that transaction performance has deteriorated on any node, check whether its transaction log is too small (see Section 25.5).

If the transaction log is big enough, but transaction performance still deteriorates, tuning the system might be necessary. *Guide to OpenVMS Performance Management* for information on tuning your system.

Example

This example shows how to monitor transaction performance on a VMScluster system that has two nodes, BLUE and RED.

Monitor transactions on nodes BLUE and RED for one day:

```
$ MONITOR TRANSACTION/SUMMARY=DISK$LOG1:[LOGFILES]TRANSACTIONS.SUM -  
_ $ /ENDING="+1-"/NODE=(BLUE,RED)
```


Managing DECdtm Services

25.4 Monitoring Transaction Performance

Examine the summary file:

```

DISTRIBUTED TRANSACTION STATISTICS
on node BLUE
SUMMARY
From: 16-MAY-1995 14:23:51
To: 17-MAY-1995 14:23:51

CUR AVE MIN MAX
Start Rate 49.02 43.21 31.30 49.02
Prepare Rate 48.70 43.23 30.67 48.70
One Phase Commit Rate 0.00 0.00 0.00 0.00
Total Commit Rate 48.70 43.19 31.30 48.70
Abort Rate 0.00 0.00 0.00 0.00
End Rate 48.70 43.19 31.30 48.70
Remote Start Rate 0.00 0.00 0.00 0.00
Remote Add Rate 0.00 0.00 0.00 0.00

Completion Rate 0-1 21.42 13.57 0.63 21.42
by Duration 1-2 25.97 29.15 24.59 33.87
in Seconds 2-3 1.29 0.47 0.00 4.47
3-4 0.00 0.00 0.00 0.00
4-5 0.00 0.00 0.00 0.00
5+ 0.00 0.00 0.00 0.00

```

SUMMARIZING

```

DISTRIBUTED TRANSACTION STATISTICS
on node RED
SUMMARY
From: 16-MAY-1995 14:23:52
To: 17-MAY-1995 14:23:52

```

Make a note of the following values:

- Average end rate.

For node BLUE, the average end rate is 43.19 transactions per second.

- Average completion rates.

For node BLUE, the average completion rates are as follows:

13.57 transactions completed in 0 to 1 seconds

29.15 transactions completed in 1 to 2 seconds

0.47 transactions completed in 2 to 3 seconds

Compare the results from this monitoring session to those of previous sessions:

Session	End Rate	Completion Rates		
		0-1 Secs	1-2 Secs	2-3 Secs
June	42.13	12.98	28.13	1.02
July	38.16	10.35	25.80	2.01
August	43.19	13.57	29.15	0.47

The results for node BLUE show no signs of deteriorating performance.

25.5 Checking Whether a Transaction Log Is Too Small

If transaction performance has deteriorated on a node, check whether its transaction log is too small.

Section 25.4 describes how to find out whether transaction performance has deteriorated.

Managing DECdtm Services

25.5 Checking Whether a Transaction Log Is Too Small

How to Perform This Task

1. Log in to the node that the transaction log belongs to.
2. Check how many times the transaction log has stalled, using LMCP's **SHOW LOG/CURRENT** command:

```
$ RUN SYSSYSTEM:LMCP
LMCP> SHOW LOG/CURRENT
```

Note the number of checkpoints and stalls displayed by this command.
3. Wait for five minutes, then repeat the **SHOW LOG/CURRENT** command.
Note the number of checkpoints and stalls again.
4. Compare the information from the **SHOW LOG/CURRENT** commands:
If the number of checkpoints has not changed, wait until the system is busier, then try this task again.
If the number of checkpoints has increased, and the number of stalls has increased by more than one, the transaction log is too small.
5. If the transaction log is too small, increase its size. For information on how to change the size of a transaction log, see Section 25.6.

Example

This example shows how to check whether node BLUE's transaction log is too small.

Log in to node BLUE. Then check how many times the transaction log has stalled:

```
$ RUN SYSSYSTEM:LMCP
LMCP> SHOW LOG/CURRENT
```

```
Checkpoint starts/ends      2464/2464
Stall starts/ends           21/21
Log status: no checkpoint in progress, no stall in progress.
```

The number of checkpoints is 2464, and the transaction log has stalled 21 times.

Wait for five minutes, then repeat the **SHOW LOG/CURRENT** command:

```
LMCP> SHOW LOG/CURRENT

Checkpoint starts/ends      2514/2514
Stall starts/ends           28/28
Log status: no checkpoint in progress, no stall in progress.
```

The number of checkpoints has increased since the previous reading, and the transaction log has now stalled 28 times, an increase of 7. This means that the transaction log is too small.

25.6 Changing the Size of a Transaction Log

To determine if changing the size of a transaction log is necessary, see Section 25.5.

How to Perform This Task

Caution

Follow all the steps carefully. Taking shortcuts can lead to data corruption.

Managing DECdtm Services

25.6 Changing the Size of a Transaction Log

1. Log in to the node that the transaction log belongs to.
2. Find out which directory the transaction log is in, using LMCP's SHOW LOG command:

SHOW LOG SYSTEM\$node.LM\$JOURNAL

where *node* is the name of the node that the transaction log belongs to.

3. Rename the transaction log:

RENAME dirspecSYSTEM\$node.LM\$JOURNAL dirspecSYSTEM\$node.LM\$OLD

where:

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.

4. Can you stop all the software that uses DECdtm services without shutting down any nodes?

Yes Close the transaction log as follows:

Step	Action
a.	Stop all the software that uses DECdtm services.
b.	Close the transaction log using LMCP's CLOSE LOG command: \$ RUN SYS\$SYSTEM:LMCP LMCP> CLOSE LOG The CLOSE LOG command closes the transaction log and stops the DECdtm TP_SERVER process. The command fails if any software is using DECdtm services.
c.	Did the CLOSE LOG command succeed? Yes Restart the TP_SERVER process: \$ @SYS\$STARTUP:DECDTM\$STARTUP.COM No Wait for 30 seconds, then repeat steps 4b and 4c.

No Close the transaction log by rebooting the node. Log in to the node when it has rebooted.

5. Change the size of the transaction log, using LMCP's CONVERT LOG command:

CONVERT LOG/SIZE=size dirspecSYSTEM\$node.LM\$OLD dirspecSYSTEM\$node.LM\$JOURNAL

where:

size is the new size of the transaction log in blocks.

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.

6. If you stopped the software that uses DECdtm services in step 4, restart the software.

7. Delete the old transaction log:

DELETE dirspecSYSTEM\$node.LM\$OLD;

where:

Managing DECdtm Services

25.6 Changing the Size of a Transaction Log

dirspec is the full specification of the directory containing the old transaction log.

node is the name of the node that the transaction log belongs to.

Example

This example shows how to change the size of node RED's transaction log to 6000 blocks. Node RED is in a VMScluster, and its transaction log is in DISK\$LOG2:[LOGFILES].

Log in to node RED. Find out which directory RED's transaction log is in, then rename the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG SYSTEM$RED.LM$JOURNAL

Directory of DISK$LOG2:[LOGFILES]

SYSTEM$RED.LM$JOURNAL;1

Total of 1 file.
LMCP> EXIT
```

```
$ RENAME DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL -
_$ DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD
```

Stop all software that uses DECdtm services. Then close the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
Transaction log closed, TP_SERVER process stopped
LMCP> EXIT
```

Restart the TP_SERVER process:

```
$ @ SYS$STARTUP:DECDTM$STARTUP.COM
```

Change the size of the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CONVERT LOG/SIZE=6000 DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD -
_LMCP> DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL
Log file DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL;1 created.
Log file DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD converted.
LMCP> EXIT
```

Restart the software that uses DECdtm services.

Delete the old transaction log:

```
$ DELETE DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD;
```

25.7 Moving a Transaction Log

You may want to move a transaction log if:

- You want to place the transaction log on a faster disk
- You want to redistribute the work load on your disks

Managing DECdtm Services

25.7 Moving a Transaction Log

How to Perform This Task

Caution

Follow all the steps carefully. Taking shortcuts can lead to data corruption.

1. Decide the location that you want to move the transaction log to, using the guidelines in Section 25.2.2. Remember that the disk must have enough contiguous space to hold the transaction log.
2. Log in to the node that the transaction log belongs to.
3. If you are in a VMSccluster environment, make sure that the disk you want to move the transaction log to is mounted clusterwide.
4. Decide which directory you want to move the transaction log to. You may want to create a new directory for the transaction log.
5. Find out which directory the transaction log is in, using LMCP's SHOW LOG command:

```
SHOW LOG SYSTEM$node.LM$JOURNAL
```

where *node* is the name of the node that the transaction log belongs to.

6. Rename the transaction log:

```
RENAME dirspecSYSTEM$node.LM$JOURNAL dirspecSYSTEM$node.LM$OLD
```

where:

dirspec is the full specification of the directory containing the transaction log.
node is the name of the node that the transaction log belongs to.

7. Can you stop all the software that uses DECdtm services without shutting down any nodes?

Yes Close the transaction log as follows:

Step	Action
a.	Stop all the software that uses DECdtm services.
b.	Close the transaction log using LMCP's CLOSE LOG command: \$ RUN SYS\$SYSTEM:LMCP LMCP> CLOSE LOG The CLOSE LOG command closes the transaction log and stops the DECdtm TP_SERVER process. The command fails if any software is using DECdtm services.
c.	Did the CLOSE LOG command succeed? Yes Restart the TP_SERVER process: \$ @SYS\$STARTUP:DECDTM\$STARTUP.COM No Wait for 30 seconds, then repeat steps 7b and 7c.

No Close the transaction log by rebooting the node. Log in to the node when it has rebooted.

Managing DECdtm Services

25.7 Moving a Transaction Log

8. Make sure that SYS\$JOURNAL points to the directory that you want to move the log to. If SYS\$JOURNAL does not point to this directory, redefine SYS\$JOURNAL:

```
DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL dirspec[,...]
```

where *dirspec* is the full specification of a directory containing one or more transaction logs. List all the directories that will contain transaction logs after you have moved the transaction log. You can list the directories in any order.

In a VMScluster environment, use SYSMAN to redefine SYS\$JOURNAL clusterwide.

9. If you redefined SYS\$JOURNAL in step 8, edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the definition of SYS\$JOURNAL.

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

10. Move the transaction log, using LMCP's CONVERT LOG command:

```
CONVERT LOG old-dirspecSYSTEM$node.LM$OLD new-dirspecSYSTEM$node.LM$JOURNAL
```

where:

<i>old-dirspec</i>	is the full specification of the directory that currently contains the transaction log.
<i>node</i>	is the name of the node that the transaction log belongs to.
<i>new-dirspec</i>	is the full specification of the directory that you are moving the transaction log to.

11. If you stopped the software that uses DECdtm services in step 7, restart the software.
12. Delete the old transaction log:

```
DELETE dirspecSYSTEM$node.LM$OLD;
```

where:

- *dirspec* is the full specification of the directory containing the old transaction log.
- *node* is the name of the node that the transaction log belongs to.

Example

This example shows how to move BLUE's transaction log. BLUE is in a VMScluster. The VMScluster members and the locations of their transaction logs are as follows:

Node	Directory Containing Log
BLUE	DISK\$LOG1:[LOGFILES]
RED	DISK\$LOG2:[LOGFILES]

Neither node has a node-specific version of SYLOGICALS.COM.

Decide where you want to move BLUE's transaction log to. In this example, assume that you want to move it to DISK\$LOG3:[LOGFILES].

Managing DECdtm Services

25.7 Moving a Transaction Log

Log in to node BLUE. Then mount the disk clusterwide, and create a new directory for the transaction log:

```
$ MOUNT/CLUSTER/SYSTEM DUA3: LOG3
$ CREATE/DIRECTORY DISK$LOG3:[LOGFILES]
```

Find out which directory BLUE's transaction log is in, then rename the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG SYSTEM$BLUE.LM$JOURNAL
```

Directory of DISK\$LOG1:[LOGFILES]

SYSTEM\$BLUE.LM\$JOURNAL;1

Total of 1 file.

```
LMCP> EXIT
```

```
$ RENAME DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL -
_$ DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD
```

Stop all software that uses DECdtm services. Then close the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
Transaction log closed, TP_SERVER process stopped
LMCP> EXIT
```

Restart the TP_SERVER process:

```
$ @SYS$STARTUP:DECDTM$STARTUP.COM
```

Redefine SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL -
SYSMAN> DISK$LOG2:[LOGFILES], DISK$LOG3:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition. Then move the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CONVERT LOG DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD -
LMCP> DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
Log file DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL;1 created.
Log file DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD converted.
LMCP> EXIT
```

Restart the software that uses DECdtm services. Then delete the old transaction log:

```
$ DELETE DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD;
```

25.8 Dismounting a Disk

Before you can dismount a disk, you must close any transaction logs on the disk. This section describes how to dismount a disk that has transaction logs.

How to Perform This Task

1. Find out which transaction logs are on the disk you want to dismount, using LMCP's SHOW LOG command:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG
```

2. Stop all the software that uses DECdtm services, if you can do so without shutting down any nodes.

If you cannot stop the software, reboot one or more nodes in step 3.

3. For each transaction log on the disk:

- a. Log in to the node that the transaction log belongs to.

- b. Rename the transaction log:

```
RENAME dirspecSYSTEM$node.LM$JOURNAL dirspecSYSTEM$node.LM$TEMP
```

where:

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.

- c. Did you stop all the software that uses DECdtm services in step 2?

Yes Close the transaction log as follows:

Step	Action
1)	Close the transaction log using LMCP's CLOSE LOG command: <pre>\$ RUN SYS\$SYSTEM:LMCP LMCP> CLOSE LOG</pre> <p>The CLOSE LOG command closes the transaction log, and stops the DECdtm TP_SERVER process. The command fails if any software is using DECdtm services.</p>
2)	Did the CLOSE LOG command succeed? <p>Yes Restart the TP_SERVER process:</p> <pre>\$ @SYS\$STARTUP:DECDTM\$STARTUP.COM</pre> <p>No Wait for 30 seconds, then repeat step 3c.</p>

No Close the transaction log by rebooting the node. When the node has rebooted, log in.

4. Dismount the disk. For instructions on how to dismount a disk, see Section 8.9

5. When you want to mount the disk again, follow these steps:

- a. Mount the disk. For instructions on how to mount a disk, see Section 8.5.

If you are in a VMScluster, mount the disk clusterwide.

- b. Rename each transaction log on the disk:

```
RENAME dirspecSYSTEM$node.LM$TEMP dirspecSYSTEM$node.LM$JOURNAL
```


Managing DECdtm Services

25.8 Dismounting a Disk

where:

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.

- c. If you stopped the software that uses DECdtm services, restart the software.

Example

This example shows how to dismount the disk DISK\$LOG3.

Find out which transaction logs are on the disk:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG
```

```
Directory of DISK$LOG3:[LOGFILES]
SYSTEM$BLUE.LM$JOURNAL;1
```

The only transaction log on DISK\$LOG3 is node BLUE's transaction log.

Stop all the software that uses DECdtm services.

Log in to node BLUE. Then rename the transaction log:

```
$ RENAME DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL -
_$ DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$TEMP
```

Close the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
Transaction log closed, TP_SERVER process stopped
LMCP> EXIT
```

Restart the TP_SERVER process:

```
$ @SYS$STARTUP:DECDTM$STARTUP.COM
```

Dismount the disk:

```
$ DISMOUNT/CLUSTER DISK$LOG3:
```

When you want to mount the disk again, mount it clusterwide:

```
$ MOUNT/CLUSTER/SYSTEM DUA3: LOG3
```

Rename BLUE's transaction log:

```
$ RENAME DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$TEMP -
_$ DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
```

Restart the software that uses DECdtm services.

25.9 Adding a Node

For every node you add to a VMSccluster, you must create a new transaction log. This section describes how to create a transaction log for a new node.

How to Perform This Task

Before you perform this task, the new node must be configured into the VMScLuster. For instructions on how to configure a node into a VMScLuster, see *VMScLuster Systems for OpenVMS*.

1. Decide the size and location of the new node's transaction log, using the guidelines in Section 25.2. Remember that the disk must have enough contiguous space to hold the log.
2. Make sure that the disk on which you want to create the transaction log is mounted clusterwide.
3. Decide which directory you want to create the new transaction log in. You may want to create a new directory for the transaction log.
4. Make sure that SYS\$JOURNAL points to the directory in which you want to create the new node's transaction log. If SYS\$JOURNAL does not point to this directory, use SYSMAN to redefine SYS\$JOURNAL clusterwide:

```
DO DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL dirspec[,...]
```

where *dirspec* is the full specification of a directory containing one or more transaction logs. List all the directories that contain transaction logs, including the directory in which you want to create the new node's transaction log. You can list the directories in any order.

5. If you redefined SYS\$JOURNAL in step 4, edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition.

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

6. Create the transaction log, using LMCP's CREATE LOG command:

```
CREATE LOG [/SIZE=size] dirspecSYSTEM$node.LM$JOURNAL
```

where:

<i>size</i>	is the size of the transaction log in blocks. By default, the size of the transaction log is 4000 blocks.
<i>dirspec</i>	is the full specification of the directory in which you want to create the transaction log.
<i>node</i>	is the name of the new node.

Example

This example shows how to create a transaction log for a new node, WHITE.

In this example, the VMScLuster members and the locations of their transaction logs are as follows:

Node	Directory Containing Log
BLUE	DISK\$LOG3:[LOGFILES]
RED	DISK\$LOG2:[LOGFILES]

Neither node has a node-specific version of SYLOGICALS.COM.

Managing DECdtm Services

25.9 Adding a Node

Decide the size and location of WHITE's transaction log:

Node	Size of Log (in Blocks)	Disk
WHITE	5000	DUA4

Mount the disk DUA4 clusterwide:

```
$ MOUNT/CLUSTER/SYSTEM DUA4: LOG4
```

Create a directory for the transaction log:

```
$ CREATE/DIRECTORY DISK$LOG4:[LOGFILES]
```

Redefine SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL -
SYSMAN> DISK$LOG2:[LOGFILES], DISK$LOG3[LOGFILES], DISK$LOG4:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$STARTUP:SYLOGICALS command procedure to update the SYS\$JOURNAL definition. Then create the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CREATE LOG/SIZE=5000 DISK$LOG4:[LOGFILES]SYSTEM$WHITE.LM$JOURNAL
LMCP> EXIT
```

25.10 Removing a Node

This section describes how to remove a node if you are using DECdtm services.

How to Perform This Task

If you have a standalone machine, perform steps 1 to 8 only.

Caution

Follow all the steps carefully. Taking shortcuts can lead to data corruption.

1. Log in to the node that you want to remove.
2. Stop all the software that uses DECdtm services.
3. Find out whether the node's transaction log contains any active transactions, using LMCP's DUMP/ACTIVE command:

```
DUMP/ACTIVE SYSTEM$node.LM$JOURNAL
```

where *node* is the name of the node that you want to remove.

This command displays details of all the active transactions. The last line gives the total number of active transactions.

4. If the transaction log contains active transactions, follow these steps:
 - a. Run recovery procedures for all software that uses DECdtm services.
 - b. Find out if the node's transaction log still contains active transactions, using LMCP's DUMP/ACTIVE command.

- c. If the transaction log still contains active transactions, contact your Customer Support Center.
5. Redefine SYS\$JOURNAL to exclude the directory that contains the transaction log of the node you want to remove, unless the directory contains other transaction logs.

DEFINE/SYSTEM/EXECUTIVE SYS\$JOURNAL dirspec[,...]

where *dirspeg* is the full specification of a directory containing one or more transaction logs. List all the directories that contain any transaction logs other than the transaction log of the node you are removing. You can list the directories in any order.

In a VMScluster, use SYSMAN to redefine SYS\$JOURNAL clusterwide.

6. If you redefined SYS\$JOURNAL in step 5, edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition.

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

7. Archive the transaction log.
8. Shut down the node.
9. Restart the software that uses DECdtm services.
10. Reconfigure the VMScluster to remove the node.

For information on how to reconfigure a VMScluster, see *VMScluster Systems for OpenVMS*.

Example

This example shows how to remove node BLUE. In this example, the VMScluster members and the locations of their transaction logs are as follows:

Node	Directory Containing Log
BLUE	DISK\$LOG3:[LOGFILES]
RED	DISK\$LOG2:[LOGFILES]
WHITE	DISK\$LOG4:[LOGFILES]

None of the nodes has a node-specific version of the SYLOGICALS.COM command procedure.

Log in to node BLUE.

Stop all the software that uses DECdtm services. Then find out if BLUE's transaction log contains any active transactions:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> DUMP/ACTIVE SYSTEM$BLUE.LM$JOURNAL
Dump of log file DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
.
.
.
Total of 0 transactions active, 0 prepared and 0 committed.
LMCP> EXIT
```


Managing DECdtm Services

25.10 Removing a Node

Redefine SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE SYS$JOURNAL -
SYSMAN> DISK$LOG2:[LOGFILES], DISK$LOG4:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition.

Archive BLUE's transaction log. Then shut down node BLUE:

```
$ @SYS$SYSTEM:SHUTDOWN.COM
```

Should an automatic system reboot be performed [NO]? NO

Restart the software that uses DECdtm services. Then reconfigure the VMScluster:

```
$ @SYS$STARTUP:CLUSTER_CONFIG.COM
```

Cluster Configuration Procedure

1. ADD a node to a cluster.
2. REMOVE a node from the cluster.
3. CHANGE a cluster member's characteristics.
4. CREATE a duplicate system disk for BLUE.

Enter choice [1]: 2

Updating network database...

The configuration procedure has completed successfully.

25.11 Disabling DECdtm Services

By default, DECdtm services start automatically when you boot the computer. The DECdtm process, TP_SERVER, then checks for a transaction log, and continues checking until it finds one.

Disable DECdtm services if you do not use, and do not plan to use, any software that uses DECdtm services. This saves memory and CPU time.

In a VMScluster, disable DECdtm services on all the nodes in the cluster.

How to Perform This Task

1. For each node:

- a. Log in to the node.
- b. Stop the TP_SERVER process using LMCP's CLOSE LOG command:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
```

The CLOSE LOG command stops the TP_SERVER process, providing no software is using DECdtm services.

If the CLOSE LOG command fails, do not continue this task.
If you have already stopped the TP_SERVER process on other nodes in a VMSccluster system, restart the process using the SYS\$STARTUP:DECDTM\$STARTUP.COM command procedure.

2. Add the following line to the SYS\$MANAGER:SYLOGICALS.COM command procedure:

```
$ !  
$ DEFINE/SYSTEM/EXECUTIVE SYS$DECDTM_INHIBIT yes  
$ !
```

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

This stops the TP_SERVER process being created the next time you boot the system.

25.12 Enabling DECdtm Services

Enable DECdtm services only if you have previously disabled them and you now want to run software that uses DECdtm services.

How to Perform This Task

1. Deassign the logical name SYS\$DECDTM_INHIBIT:

```
$ DEASSIGN/SYSTEM/EXECUTIVE SYS$DECDTM_INHIBIT
```

In a VMSccluster environment, use SYSMAN to deassign SYS\$DECDTM_INHIBIT clusterwide.

2. Start up the DECdtm services process, TP_SERVER:

```
$ @SYS$STARTUP:DECDTM$STARTUP.COM
```

In a VMSccluster environment, use SYSMAN to start up the TP_SERVER process clusterwide.

3. Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to delete the SYS\$DECDTM_INHIBIT definition. This ensures that DECdtm services start automatically when you boot the system.

Example

This example shows how to enable DECdtm services in a VMSccluster environment.

Deassign SYS\$DECDTM_INHIBIT, then start up the TP_SERVER process.

```
$ RUN SYS$SYSTEM:SYSMAN  
SYSMAN> SET ENVIRONMENT/CLUSTER  
SYSMAN> DO DEASSIGN/SYSTEM/EXECUTIVE SYS$DECDTM_INHIBIT  
SYSMAN> DO @SYS$STARTUP.DECDTM$STARTUP.COM  
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to delete the SYS\$DECDTM_INHIBIT definition.

25.13 Using DECdtm Services in a DECnet/OSI Network

If your DECdtm transactions span different VMScluster systems or standalone computers in a DECnet/OSI network, make sure the value of the SCSNODE system parameter is unique for each computer. The SCSNODE system parameter defines the name of the computer; a name must not be duplicated in a transaction group or the DECdtm transaction can fail.

25.13.1 Understanding the Configuration of a Transaction Group


Figure 25-2 shows an example of a transaction group. A transaction group conforms to the following guidelines:

- A computer can belong to only one transaction group.
- Every computer in a VMScluster system belongs to the same transaction group.
- Computers A and B belong to the same transaction group if any transaction on computer A involves computer B.
- Computers A and C belong to the same transaction group if any transaction on computer A involves computer B, and any transaction on computer B, or any node in B's VMScluster system, involves computer C.

Managing DECdtm Services

25.13 Using DECdtm Services in a DECnet/OSI Network

Figure 25-2 Transaction Group



Key:



computer



transaction

ZK-6302A-GE

In this example, transaction group members are computer TOM and all computers in clusters FRED and BILL because:

- Transactions on a computer in cluster FRED involve other computers in cluster FRED, and a computer in cluster BILL.
- Transactions on a computer in cluster BILL involve standalone system TOM.
- No other computers in the network are involved in transactions with computers in clusters FRED or BILL, or with standalone computer TOM.

Managing DECdtm Services

25.13 Using DECdtm Services in a DECnet/OSI Network

25.13.2 Determining SCSNODE Name Uniqueness

Each computer in the transaction group must have a unique SCSNODE name. Use the following guidelines to make sure the SCSNODE name is not duplicated on other computers in the transaction group. To determine SCSNODE values see Section 23.5.3.

1. Note which computers belong to the same transaction group.
2. Note the SCSNODE value for each computer in the transaction group. The value must be different from:
 - The SCSNODE values of other computers in the transaction group
 - DECnet synonyms of other computers in the entire network
 - DECnet simple names of other computers on the same local root

For information about DECnet synonyms and DECnet simple names, see the *DECnet/OSI DECdns Management* manual.
3. If a computer in the transaction group belongs to a VMSccluster, note that computer's SCSNODE value. The value must be different from:
 - DECnet simple names of other computers in the same VMSccluster
 - DECnet simple names of computers on the same local root as other VMSccluster members
4. Change any duplicate SCSNODE names. See Section 23.5.3 to assign new names.

Managing Special Processing Environments

The OpenVMS operating system supports the following special environments:

- Symmetric multiprocessing
- Vector processing (available only on certain CPU models)

This chapter describes how to set up and manage these special processing environments.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Creating a multiprocessing environment	Section 26.2.1
Monitoring a multiprocessing environment	Section 26.2.2
†Loading the vector processing support code	Section 26.4.1
†Configuring a vector processing system	Section 26.4.2
†Managing vector processes	Section 26.4.3
†Restricting access to the vector processor with ACLs	Section 26.4.4
†Obtaining information about a vector processing system	Section 26.4.5
†Loading the VAX Vector Instruction Emulation facility (VVIEF)	Section 26.4.6
†VAX specific	

This chapter explains the following concepts:

Concept	Section
Symmetric multiprocessing	Section 26.1
Primary and secondary processors	Section 26.1.1
Available and active sets	Section 26.1.2
Vector processing	Section 26.3
†VAX support for vector processing	Section 26.3.1
†The VAX Vector Instruction Emulation facility (VVIEF)	Section 26.3.2
†VAX specific	

26.1 Understanding Multiprocessing

A multiprocessing system consists of two or more CPUs that address a common pool of memory and that are capable of executing instructions simultaneously.

The OpenVMS operating system supports a tightly coupled, symmetric multiprocessing (SMP) system. In a tightly coupled SMP system, all processors execute a single copy of the operating system and have equal access to all operating system code and system resources. OpenVMS SMP dynamically selects the CPU where a process will run based on process priority.

A multiprocessing system can function as an isolated entity, a node in a network, or a member of a VMScluster environment. Multiprocessing and uniprocessing systems run the same operating system, although multiprocessing can be enabled only on selected VAX and Alpha processors. All processors in a multiprocessing environment must be at the same hardware and firmware level to guarantee that a given processor is capable of resuming the execution thread of a process that had been executing previously on another processor in the system.

26.1.1 Primary and Secondary Processors

In a multiprocessing system, one processor has the responsibility of starting other processors in the system. The **primary processor** is that processor in the system that is either logically or physically attached to the console device. As such, it is the processor that is the target of the console commands that boot the multiprocessing system. In this role, only the primary processor performs the initialization activities that define the operating system environment and prepare memory for the entire system. In addition, the primary processor serves as the system timekeeper, maintaining the system time and monitoring the timer queue for the expiration of its elements. In this sense, all processors in a multiprocessing system that do *not* have these responsibilities are known as **secondary processors**.

26.1.2 Available and Active Sets

An **available set** is made up of the processors that have passed the system's power-on hardware diagnostics and may or may not be actively involved in the system. Together, the primary and the secondary processors comprise the multiprocessing system's available set.

The **active set** is the subset of the VAX or Alpha system's processors that have passed power-on diagnostics and are actively participating in system operations. The operating system identifies each processor in these sets by its **CPU ID**, a value prevalent in the syntax and displays of certain DCL and utility commands.

26.1.3 Processor Capabilities

The processors in a multiprocessing system offer certain capabilities to the processes executing in the system. The following capabilities are supported:

- Primary
- Quorum
- Run
- Vector (VAX Only)

In addition, mechanisms exist to add and subtract other capabilities.

The Run capability affects CPU starting and stopping operations.

26.2 Managing Symmetric Multiprocessing (SMP) Environments

Managing multiprocessing systems involves creating and monitoring a multiprocessing environment.

26.2.1 Creating a Multiprocessing Environment

You can control the membership and character of a multiprocessing system at boot time by setting system parameters designed for these purposes. Among the system parameters that manage a multiprocessing system are the following:

Parameter	Function
MULTIPROCESSING	Determines which synchronization image is loaded into the operating system at boot time
SMP_CPUS	Determines which processors are brought into the multiprocessing environment from the available set at boot time

For more information about these and other system parameters, see the *OpenVMS System Management Utilities Reference Manual*.

You can add an available processor to the active set at boot time, or you can add it later using the DCL command START/CPU. The DCL command STOP/CPU removes a processor from the active set.

Symmetric Multiprocessing Extension License

Alpha

On Alpha systems, you must register the SMP Extension License if you have an SMP system. This license upgrades the Operating System Base License and all Interactive User licenses to the matching multiprocessing level of your system.

Because the SMP Extension License grants all the rights the existing Base and User licenses provide at the uniprocessing level, reinstalling those licenses when you upgrade to a multiprocessing system is unnecessary. When your system is upgraded to a new multiprocessing level, add an SMP Extension License to your existing license. ♦

26.2.2 Monitoring a Multiprocessing Environment

Several operating system features provide special information about the character, capabilities, and status of a multiprocessor system. They include the DCL command SHOW CPU and the Monitor utility.

Obtaining Information About a Multiprocessor Configuration

The SHOW CPU command displays three levels of information describing the configuration and status of a multiprocessing system:

Level	Command Example	Display Contents
Summary	SHOW CPU	Indicates which processor is primary, which processors are configured, and which processors are active; displays the minimum revision levels for processors in the system and the setting of the MULTIPROCESSING system parameter; and indicates whether multiprocessing is enabled.

Managing Special Processing Environments

26.2 Managing Symmetric Multiprocessing (SMP) Environments

Level	Command Example	Display Contents
Brief	SHOW CPU/BRIEF	Produces information from the summary display; lists the current CPU state and the current process (if any) for each configured processor.
Full	SHOW CPU/FULL	Produces information from the summary display and displays additional information; lists the current CPU state, current process (if any), revision levels, and capabilities for each configured processor; indicates which processes can be executed only on certain processors.

For more information about the DCL commands relating to SMP, see the *OpenVMS DCL Dictionary*; for information about the Monitor utility, see the MONITOR section in the *OpenVMS System Management Utilities Reference Manual*.

26.3 Understanding Vector Processing

A single data item, having one value, is known as a **scalar**. A group of related scalar values, or elements, all of the same data type, is known as a **vector**.

Traditional (scalar) computers operate only on scalar values, and must process vector elements sequentially. Vector computers, on the other hand, recognize vectors as native data structures and can operate on an entire vector with a single vector instruction. Because this type of processing involves the concurrent execution of multiple arithmetic or logical operations, a vector computer can routinely process a vector four to five times faster than a traditional computer can using only scalar instructions.

Vector processors gain a further speed advantage over scalar processors by their use of special hardware techniques designed for the fast processing of streams of data. These techniques include data pipelining, chaining, and other forms of hardware parallelism in memory and in arithmetic and logical functional units. Pipelined functional units allow the vector processor to overlap the execution of successive computations with previous computations.

26.3.1 VAX Support for Vector Processing (VAX Only)

VAX

The VAX vector architecture includes sixteen 64-bit vector registers (V0 through V15), each containing 64 elements; vector control registers, including the vector count register (VCR), vector length register (VLR), and vector mask register (VMR); vector functional units; and a set of vector instructions. VAX vector instructions transfer data between the vector registers and memory, perform integer and floating-point arithmetic, and execute processor control functions. A more detailed description of the VAX vector architecture, vector registers, and vector instructions appears in the *VAX MACRO and Instruction Set Reference Manual*.

Those VAX systems that comply with the VAX vector architecture are known as **vector-capable** systems.

A VAX vector processing system configuration includes one or more integrated scalar-vector processor pairs, or **vector-present processors**. Such a configuration can be symmetric, including a vector coprocessor for each scalar, or asymmetric, incorporating additional scalar-only processors. Depending upon the model of the VAX vector processing system, the scalar and vector CPUs of vector-present processors can be either a single, integral physical module or separate, physically independent modules. In either case the scalar and vector

Managing Special Processing Environments

26.3 Understanding Vector Processing

CPUs are logically integrated, sharing the same memory and transferring data over a dedicated, high-speed internal path.

Like VAX scalar processing systems, a VAX vector processing system can participate as a member of a VAXcluster or a node in a network, or be run as a standalone system. ♦

26.3.2 VAX Vector Instruction Emulation Facility (VVIEF) (VAX Only)

The VAX Vector Instruction Emulation facility (VVIEF) is a standard feature of the OpenVMS operating system that allows vectorized applications to be written and debugged in a VAX system in which vector processors are not available. VVIEF emulates the VAX vector processing environment, including the nonprivileged VAX vector instructions and the vector system services. Use of VVIEF is restricted to user mode code.

VVIEF is strictly a program development tool, and *not* a run-time replacement for vector hardware. Vectorizing applications to run under VVIEF offers no performance benefit; vectorized applications running under VVIEF execute more slowly than their scalar counterparts.

The operating system supplies the VVIEF bootstrap code as an executive loadable image. Note that, in the presence of OpenVMS vector support code, VVIEF remains inactive. Although it is possible to prevent the loading of vector support code in a vector-present system (see Section 26.4.1) and activate VVIEF, there are few benefits.

See Section 26.4.6 for additional information on loading and unloading VVIEF. ♦

26.4 Managing the Vector Processing Environment (VAX Only)

VAX

The following sections describe tasks for managing a vector processing system.

26.4.1 Loading the Vector Processing Support Code (VAX Only)

By default, in a VAX vector processing system, the system automatically loads the vector processing support code at boot time. You can override the default behavior by setting the static system parameter VECTOR_PROC as described in Table 26-1.

Table 26-1 Settings of VECTOR_PROC System Parameter (VAX Only)

Value	Result
0	Do not load the vector processing support code, regardless of the system configuration.
1	Load the vector processing support code if at least one vector-present processor exists. This is the default value.
2	Load the vector processing support code if the system is vector-capable. This setting is most useful for a system in which processors have separate power supplies. With this setting, you can reconfigure a vector processor into the system without rebooting the operating system.

Managing Special Processing Environments

26.4 Managing the Vector Processing Environment (VAX Only)

26.4.2 Configuring a Vector Processing System (VAX Only)

You can add a vector-present processor to or remove it from a multiprocessing configuration at boot time by using the system parameter `SMP_CPUS`, or at runtime by using the DCL commands `START/CPU` and `STOP/CPU`. Note that the operating system treats the scalar and vector CPU components of a vector-present processor as a single processor, starting them and stopping them together.

At boot time, the setting of the system parameter `SMP_CPUS` identifies which secondary processors in a multiprocessing system are to be configured, including those processors that are vector present. (the operating system always configures the primary processor.) The default value of `-1` boots all available processors, scalar and vector-present alike, into the configuration. (See the *OpenVMS System Management Utilities Reference Manual* for additional information on this parameter.) Note that, prior to starting a vector-present processor, you should ensure that the vector processing support code (see Section 26.4.1) is loaded at boot time. Otherwise, processes will be able to use only the scalar CPU component of the vector-present processor.

To bring secondary processors into a running multiprocessing system, you use the DCL command `START/CPU`. To remove secondary processors from the system, use the `STOP/CPU` commands. Again, you must ensure that the vector processing support code has been loaded at boot time for the vector CPU component of vector-present processors started in this way to be used.

Note, however, that a `STOP/CPU` command fails and generates a message if it would result in the removal of a vector-present processor that is the sole provider of the vector capability for currently active vector consumers. In extreme cases, such as the removal of a processor for repair, you can override this behavior by issuing the command `STOP/CPU/OVERRIDE`. This command stops the processor, despite stranding processes.

When a `STOP/CPU/OVERRIDE` command is issued for a vector-present processor, or when a vector-present processor fails, the operating system puts all stranded vector consumers into a CPU-capability-wait (`RSN$_CPUCAP`) state until a vector-present processor is returned to the configuration. To any other process that subsequently issue a vector instruction (including a marginal vector consumer), the system returns a "requested CPU not active" message (`CPUNOTACT`).

See the *OpenVMS DCL Dictionary* for additional information on the `START/CPU` and `STOP/CPU` commands.

26.4.3 Managing Vector Processes (VAX Only)

The operating system scheduling algorithms automatically distribute vector and scalar processing resources among vector consumers, marginal vector consumers, and scalar consumers. However, VAX vector processing configurations vary in two important ways:

- The amount of vector processing activity the configuration must accommodate
- The number of vector-present processors that are available in the configuration to service vector processing needs

In a configuration that has more vector consumers in a system than scalar-vector processor pairs to service them, vector consumers share vector-present processors according to process priority. At a given priority, the system schedules vector consumers on a vector-present processor in a round-robin fashion. Each time the system must schedule a new vector consumer on a vector-present processor,

Managing Special Processing Environments

26.4 Managing the Vector Processing Environment (VAX Only)

it must save the vector context of the current vector consumer in memory and restore the vector context of the new vector consumer from memory. When such "slow" vector context switches occur too frequently, a significant portion of the processing time is spent on vector context switches relative to actual computation.

Systems that have heavy vector processing needs should be adequately configured to accommodate those needs. However, some mechanisms are available for tuning the performance of an existing configuration.

26.4.3.1 Adjusting System Resources and Process Quotas (VAX Only)

Systems in which several vector consumers are active simultaneously may experience increased paging activity as processes share the available memory. To reduce process paging, you may need to use the Authorize utility to adjust the working set limits and quotas of the processes running vectorized applications. (See the AUTHORIZE section of the *OpenVMS System Management Utilities Reference Manual* for additional information.) An increase of the process maximum working set size (system parameter WSMAX) may also be necessary. Additionally, a vectorized application may use the Lock Pages in Working Set system service (SYS\$LKWSET) to enhance its own performance.

The system allots to each vector consumer 8 KB of system nonpaged dynamic memory in which the operating system stores vector context information. Depending upon how many vector consumers may be active in the system simultaneously, you may need to adjust the system parameter NPAGEDYN. The DCL command SHOW MEMORY/POOL/FULL displays the current size of nonpaged pool in bytes.

To obtain optimal performance of a VAX vector processing system, you should take some care in setting up generic batch queues that avoid saturating the system's vector resources. If a queue contains more active vectorized batch jobs than vector-present processors in the system, a significant portion of the processing time will be spent on vector context switches.

The recommended means for dispatching vectorized batch jobs to a VAX vector processing system is to set up a separate queue (for instance, VECTOR_BATCH) with a job limit equal to the number of vector-present processors in the system. When submitting vectorized batch jobs, users should be encouraged to submit them to this generic vector-processing batch queue.

26.4.3.2 Distributing Scalar and Vector Resources Among Processes (VAX Only)

As a vector consumer, a process must be scheduled only on a vector-present processor. If the image the process is executing issues only scalar instructions for a period of time, and it must share the scalar-vector processor pair with other vector consumers, its inability to run on an available scalar processor could hamper its performance and the overall performance of the system.

By default, the operating system assumes that if a vector consumer has not issued a vector instruction for a certain period of time, it is unlikely that it will issue a vector instruction in the near future. The system relinquishes this process's need for the vector capability, classifying it as a marginal vector consumer.

In an asymmetric vector-processing configuration, detection of marginal vector consumers achieves the following desirable effects:

- Because a marginal vector consumer is eligible to run on a larger set of processors, its response time will improve.
- The scheduling of marginal vector consumers on scalar processors reduces the contention for vector-present processors.

Managing Special Processing Environments

26.4 Managing the Vector Processing Environment (VAX Only)

- Because vector consumers issuing vector instructions are more likely to be scheduled on vector-present processors, the vector CPU is more efficiently used.

Use the VECTOR_MARGIN system parameter to establish the interval of time at which the system checks the status of all vector consumers. The VECTOR_MARGIN parameter accepts an integer value between 1 and FFFFFFFF₁₆. This value represents a number of consecutive process quanta (as determined by the system parameter QUANTUM). If the process has not issued any vector instructions in the specified number of quanta, the system declares it a marginal vector consumer.

The default value of the VECTOR_MARGIN parameter is 200₁₀.

26.4.4 Restricting Access to the Vector Processor by Using ACLs (VAX Only)

A vector **capability** is a software abstract by which the operating system makes the services of the vector processor available to users. You can restrict the use of the vector processor to users holding a particular identifier by associating an access control list (ACL) with the vector capability object.

For example, a university might limit use of the vector processor to faculty and students in an image processing course, or a service bureau might charge users for access to the vector capability, time spent on the vector processor, or both.

Use the DCL command SET ACL in the following format to establish access control entries (ACEs) on a vector capability:

```
SET ACL/OBJECT=CAPABILITY VECTOR/ACL[=(ace[,...])]
```

Note that you must be in the SYSTEM user category (as described in the *OpenVMS User's Manual*) to set an ACL on the vector capability.

The following DCL command displays the ACL on the vector capability:

```
$ SHOW ACL/OBJECT=CAPABILITY VECTOR
```

Note that the ACL is on the vector capability, not on the use of any or all vector-present processors in the system. The operating system will still schedule processes without permission to use the vector capability on a vector-present processor. However, these processors will be able to use only the scalar CPU component of the processor, and cannot execute vector instructions. Likewise, because the ACL is on the vector capability and not on a vector-present processor, you cannot establish an ACL to force long-running jobs to a specific processor.

For additional information on the SET ACL and SHOW ACL commands, see the *OpenVMS DCL Dictionary*.

26.4.5 Obtaining Information About a Vector Processing System (VAX Only)

You can obtain information about the status of the vector processing system and the use of the system by individual processes through various means, including:

- The DCL lexical functions F\$GETJPI and F\$GETSYI
- The DCL command SHOW CPU
- The DCL commands SHOW PROCESS and LOGOUT/FULL
- The Accounting utility
- The Monitor utility

Managing Special Processing Environments

26.4 Managing the Vector Processing Environment (VAX Only)

26.4.5.1 DCL Lexical Functions F\$GETJPI and F\$GETSYI (VAX Only)

The DCL lexical function F\$GETJPI accepts the following items and returns the corresponding information regarding the vector status of a specified process:

Item	Return Type	Information Returned
FAST_VP_SWITCH	Integer	Number of times this process has issued a vector instruction that resulted in an inactive vector processor being enabled without the expense of a vector context switch
SLOW_VP_SWITCH	Integer	Number of times this process has issued a vector instruction that resulted in an inactive vector processor being enabled with a full vector context switch
VP_CONSUMER	Boolean	Flag indicating whether the process is a vector consumer
VP_CPUTIM	Integer	Total amount of time the process has accumulated as a vector consumer

The DCL lexical function F\$GETSYI accepts the following items and returns the corresponding information regarding the status of the vector processing system:

Item	Return Type	Information Returned
VECTOR_EMULATOR	Integer	Flag indicating the presence of the VAX Vector Instruction Emulation facility (VVIEF) in the system
VP_MASK	Integer	Mask indicating which processors in the system have vector coprocessor
VP_NUMBER	Integer	Number of vector processors in the system

See the *OpenVMS DCL Dictionary* for additional information about the DCL lexicals F\$GETJPI and F\$GETSYI.

26.4.5.2 SHOW CPU/FULL Command (VAX Only)

The SHOW CPU/FULL command lists the capabilities of the specified CPU. Issue this command to determine the presence of the vector capability in the system prior to executing a STOP/CPU command.

See the *OpenVMS DCL Dictionary* for additional information about the SHOW CPU command.

26.4.5.3 SHOW PROCESS and LOGOUT/FULL Commands (VAX Only)

If the target process has accrued any time as a vector consumer scheduled on a vector-present processor, the DCL commands SHOW PROCESS and LOGOUT /FULL display the elapsed vector CPU time and the charged vector CPU time, respectively.

To accumulate vector CPU time, a process must be a vector consumer (that is, require the system vector capability) and be scheduled on a vector-present processor. The operating system still charges the vector consumer vector CPU time, even if, when scheduled on the vector-present processor, it does not actually use the vector CPU. Note that, because scalar consumers and marginal vector consumers do not use the vector CPU, they do not accrue vector CPU time, even when scheduled on a vector-present processor.

See the *OpenVMS DCL Dictionary* for additional information about the SHOW PROCESS and LOGOUT commands.

Managing Special Processing Environments

26.4 Managing the Vector Processing Environment (VAX Only)

26.4.6 Loading the VAX Vector Instruction Emulation Facility (VVIEF) (VAX Only)

The VAX Vector Instruction Emulation facility (VVIEF) is a standard operating system feature that allows vectorized applications to be written and debugged in a VAX system in which vector processors are not available. VVIEF is intended strictly as a program development tool, and *not* as a run-time replacement for vector hardware. Vectorizing applications to run under VVIEF offers no performance benefit; vectorized applications running under VVIEF will execute more slowly than their scalar counterparts.

To cause the system to load VVIEF at the *next* system boot and at each subsequent system boot, invoke the command procedure `SYSS$UPDATE:VVIEF$INSTAL.COM`. To unload VVIEF, invoke the command procedure `SYSS$UPDATE:VVIEF$DEINSTAL.COM` and reboot the system.

You can determine the presence or absence of VVIEF in a system by issuing the following DCL commands:

```
$ X = F$GETSYI("VECTOR_EMULATOR")
$ SHOW SYMBOL X
X = 1   Hex = 00000001   Octal = 0000000001
```

A return value of 1 indicates the presence of VVIEF; a value of 0 indicates its absence.

Note that, although VVIEF may be loaded into the system, in the presence of vector support code, it remains inactive. Although it is possible to prevent the loading of vector processing support code in a vector-present system (see Section 26.4.1) and activate VVIEF, there are few benefits. Should the only vector-present processor in the system fail, the execution of preempted vectorized applications will not resume under VVIEF. ♦

Files-11 Disk Structure

This appendix explains disk terminology and disk concepts. It also describes reserved files, points out those files used by the Analyze/Disk_Structure utility (ANALYZE/DISK_STRUCTURE), and compares Files-11 On-Disk Structure (ODS) Level 1 and Files-11 ODS Level 2.

A.1 Disk Concepts

This section defines terms related to both the physical and the logical organization of disks.

A.1.1 Logical Organization of a Disk

The smallest addressable unit of information on a disk is a **block**. Files-11 On-Disk Structures define a block to consist of 512 8-bit bytes. Blocks can be treated as units for transfer between a Files-11 disk volume and memory. Files-11 ODS, however, views a disk as an array of blocks, and is generally not concerned with individual blocks.

Blocks are logically grouped into **clusters**, which are the basic units by which disk space is allocated. You determine the number of blocks in a cluster when a given disk, known as a **volume**, is first prepared for use (initialized). Cluster sizes vary for different media types. The smaller cluster sizes in the range are usually more practical. In general, a disk with a relatively small number of blocks is given a smaller cluster size, while larger disks are given larger cluster sizes to minimize the overhead for disk space allocation.


Contiguous clusters allocated to a particular file are called **extents**. An extent can contain all or part of a file. If enough contiguous area is available on the disk, the entire file is allocated as a single extent. Sometimes, however, not enough contiguous area is available to hold the entire file, or, when you create a file initially, you might not want to reserve the entire required amount of space. When the file is eventually extended, it is unlikely that the adjacent clusters will still be unallocated. If the adjacent clusters are already allocated to another file, the extension does not occur contiguously.

If a file is divided into two or more parts, each part is an extent. Thus, a file can consist of multiple extents located in separate areas on the disk, as shown in Figure A-1. Note that the file extensions are done automatically.

Files-11 Disk Structure

A.1 Disk Concepts

Figure A-1 File Extents



ZK-0738-GE


A.1.2 Physical Organization of a Disk

The smallest unit discernible to the Files-11 structure is the **sector**; for most Files-11 disks, a sector is equivalent to a block, which is 512 bytes. Other basic terms related to disks are **track** and **cylinder**. A track is the collection of sectors (or blocks, on Files-11 structures) at a single radius on one recording surface of a disk. It is accessible to a given read/write head position on the disk device. A cylinder consists of all tracks at the same radius on all recording surfaces of a disk.

Because access to any of the blocks in a given cylinder does not require any movement of the disk's read/write heads, it is generally advantageous to keep related data blocks in the same cylinder. For this reason, when choosing a cluster size for a large-capacity disk, you should usually select a cluster size that divides evenly into the cylinder size.

Figure A-2 is a graphic representation of disk tracks and cylinders.

Figure A-2 Tracks and Cylinders



ZK-0740-GE

A.2 Files-11 Structure

The Files-11 structure creates a set of nondeletable reserved files when a volume or volume set is initialized. These files control the organization of a Files-11 disk. A Files-11 structure resides on a volume, which is a physical medium such as a disk pack. A Files-11 volume is an ordered set of 512-byte blocks. The blocks are numbered consecutively from 0 to $n-1$; the value of $n-1$ is the size of the disk in blocks.

Files-11 Disk Structure

A.2 Files-11 Structure

A.2.1 File Identification (FID)

Each file on a Files-11 disk is identified by a unique, system-assigned file identification (FID) and can have a user-assigned alphanumeric name. The primary function of a Files-11 directory is to associate the user-assigned alphanumeric name of each file with the unique FID of the file. This association ensures that files present on a volume are retrievable by name.

The FID of a file consists of a set of three numbers. The first is the **file number** (NUM). The file system uses this number as an offset into the index file (reserved file INDEXF.SYS), which stores information for all files on a volume.

The second part of the FID is the **file sequence number** (SEQ), which represents the number of times a particular file number has been used. File numbers are allocated and deallocated as files are created and deleted. As a result, the file number alone cannot uniquely identify the file. By incrementing the sequence number each time a file number is used, the file system ensures that each file has a unique identification in INDEXF.SYS.

The third number in the FID is the **relative volume number** (RVN). This number indicates the volume (of a volume set) on which the file resides (ODS-2 only). If the volume set consists of a single volume, the RVN of all files on that volume is 1.

A.2.2 ODS Directory Hierarchies

The Files-11 ODS-2 structure is a multilevel directory hierarchy. The top level of the directory structure is the master file directory (MFD). The MFD of a volume is always named [000000]. The MFD contains all top-level directories, including itself, and reserved files.

A directory is a file that contains other files. A file contained in a directory can also be a directory and contain other files. By nesting directories, users can construct directory hierarchies up to nine levels deep (including the master file directory).

In a volume set, the MFD for all of the user directories on the volume set is located on relative volume 1. The entries of this MFD point to directories located on any volume in the set. These directories in turn point to files and subdirectories on any volume in the set. The MFD of any remaining volume in the set includes only the names of the reserved files for that volume.

VAX

On VAX systems, the Files-11 ODS-1 structure supports a two-level directory hierarchy. Each user identification code (UIC) is associated with a user file directory (UFD). Each UFD is included in the master file directory (MFD) of the volume. ♦

A.3 Reserved Files

This section describes the reserved files that Files-11 uses. Note that all reserved files have constant FIDs.

This section also points out the files ANALYZE/DISK_STRUCTURE uses. ANALYZE/DISK_STRUCTURE makes an in-memory copy of what these files should look like and compares it with the current version. The utility reports and repairs (if you specify the /REPAIR qualifier) any discrepancies found during these comparisons.

Files-11 Disk Structure

A.3 Reserved Files

Table A-1 shows the reserved files used by Files-11 Level 1 and Level 2, and files used by ANALYZE/DISK_STRUCTURE.

Table A-1 Reserved Files

Reserved File	File Name	†Structure Level 1	Structure Level 2	ANALYZE/DISK_STRUCTURE
Index file	INDEXF.SYS;1	X	X	X
Storage bit map file	BITMAP.SYS;1	X	X	X
Bad block file	BADBLK.SYS;1	X	X	
Master file directory	000000.DIR;1	X	X	X
Core image file	CORIMG.SYS;1	X	X	
Volume set list file	VOLSET.SYS;1		X	X
Continuation file	CONTIN.SYS;1		X	
Backup log file	BACKUP.SYS;1		X	
Pending bad block	BADLOG.SYS;1		X	
Quota file	QUOTA.SYS			X
Volume security profile	SECURITY.SYS		X	
†VAX specific				

A.3.1 Index File, INDEXF.SYS

Every Files-11 volume has an index file, which is created when the volume is initialized. (You cannot use a disk as a Files-11 disk until it has been initialized with the INITIALIZE command.)

INDEXF.SYS is a large, extendable file made up of several sections. These sections provide the operating system with the information necessary to identify a Files-11 volume, initially access that volume, and locate all the files on that volume (including INDEXF.SYS itself).

Table A-2 shows the information that is in INDEXF.SYS. Following the table are additional explanations of boot block, home block, and file headers.

Table A-2 Contents of Files-11 Index File

Term	Definition
Boot block	Virtual block number 1 of the index file. If the volume is a system volume, the boot block contains a boot program that loads the operating system into memory. If the volume is not a system volume, the boot block contains a program that displays the message that the volume is not the system device but a device that contains users' files only.
Home block	Establishes the specific identity of the volume, providing such information as the volume name and protection, the maximum number of files allowed on the volume, and the volume ownership information. The home block is virtual block number 2 of the index file.
Backup home block	A copy of the home block; permits the volume to be used even if the primary home block is destroyed.

(continued on next page)

Files-11 Disk Structure

A.3 Reserved Files

Table A-2 (Cont.) Contents of Files-11 Index File

Term	Definition
Backup index file header	Permits data on the volume to be recovered if the index file header is corrupted; occupies virtual blocks $v * 3 + 1$ through $v * 4$, where v is the volume cluster factor.
Index file bitmap	Controls the allocation of file headers and thus the number of files on the volume; contains a bit for each file header allowed on the volume. If the value of a bit for a given file header is 0, a file can be created with this file header. If the value is 1, the file header is already in use.
File headers	Makes up the bulk of the index file; contain all the information needed for gaining access to the file. Each file header describes one file on the volume. A file header contains information such as the owner UIC, protection code, creation date and time, and access control lists (ACLs); it also contains a list of extents that make up the file, describing where the file is logically located on the volume. Note that a file header can also be an extension header.
Alternate index file header	Permits recovery of data on the volume if the primary index file header becomes damaged.

A.3.1.1 Boot Block

Block 0 on a system disk is the **boot block**. It contains the location and size of the **bootstrap image**, which is used to boot the system. Certain processors, in order to boot, must read this boot block to obtain the location of the bootstrap image. For more details, see Section 4.6.

A.3.1.2 Home Block

The **home block** is normally the next block after the boot block; it identifies the disk as a Files-11 volume. If for some reason the home block cannot be read (physically unusable), an alternative block will be selected for use as the home block. This block provides specific information about the volume and default values for files on the volume. Among the items in the home block are the following:

- The volume name
- Information to locate the remainder of the index file
- The maximum number of files that can be present on the volume at any one time
- The user identification code (UIC) of the owner of the volume
- Volume protection information (specifies which users can read or write the entire volume)

Files-11 volumes contain several copies of the home block to ensure against accidental destruction of this information and the consequent loss of access to files on the volume.

A.3.1.3 File Headers

Most of the index file consists of **file headers**; each file header describes a portion of a file on the volume. File headers contain information such as the owner UIC, protection code, creation date and time, and access control lists (ACLs). Most importantly, the file header contains a list of extents that make up the file, describing where the file is logically located on the volume. If a file has a large number of extents, multiple file headers may be used to describe them. A file identifier number is associated with each file header.

When you create a file, you normally specify a file name to OpenVMS RMS, which assigns this name to the file on a Files-11 volume. OpenVMS RMS places the file name and file identifier associated with the newly created file into a directory, which contains an entry defining the location for each file. When you access the file, you supply the file name, which supplies a path to the file identifier through the directory entry. The file identifier, in turn, points to the location of the file header, which contains a listing of the extent or extents that locate the actual data.

Because they represent the current state of file storage on a volume, file headers are of particular interest to ANALYZE/DISK_STRUCTURE. Each file on a Files-11 disk (INDEXF.SYS included) is identified and located by a primary header (and extension headers, if required) in INDEXF.SYS.

Each fixed-length header contains both constant and variable-length data. This data is stored in one of the six areas shown in Table A-3.

Table A-3 Areas of Data in File Headers

Area of Data	Description
Header	This area contains the header identification, the file number and its sequence number, the protection code for the file, and offsets to the other file header areas.
Ident	This area contains the identification and accounting data for the file (for example, the name of the file, its creation date and time, and backup date and time).
Map	This area contains a list of retrieval pointers that map the virtual blocks of the file to the logical blocks of the volume. Each pointer describes one group of consecutively numbered logical blocks that is allocated to the file. Retrieval pointers are arranged in the order of the virtual blocks they represent.
Access control list	An optional area that contains ACL-related information.
Reserved	This area is reserved for use by special applications.
End checksum	The last two bytes of the file header contain a 16-bit additive checksum of the preceding 255 words of the file header. The checksum helps verify that the block is a valid file header.

A set of contiguous clusters is known as an **extent**. The size of an extent varies according to the number of contiguous clusters. For example, assume a file requires 1000 blocks of storage, and the file system finds a set of 800 contiguous blocks and a set of 200 contiguous blocks. The file would then be stored in two extents: one consisting of 800 blocks, the other of 200.

Files-11 Disk Structure

A.3 Reserved Files

The **primary header** of a file points to the first extent of that file and to as many extents as can be stored in the map area of the primary header. When the number of extents required to contain a file exceeds the map area available in the primary header, or the ACL is too large to fit in the primary header, the file is allocated an **extension header**. Extension headers contain all the constant data of the primary header, as well as the variable data (in the header map area and access control list) that specifies the locations of the extents to which the extension header points.

ANALYZE/DISK_STRUCTURE confirms the validity of a file by working its way down the list of primary and extension headers of the file. During this process, ANALYZE/DISK_STRUCTURE checks the validity of the file header, the chain of pointers to all extension headers, the retrieval pointers in all headers, and the attributes of the file.

A.3.2 Storage Bit Map File, BITMAP.SYS

The storage bit map file is a contiguous file that the file system uses to keep track of the available space on a volume. This file contains a storage control block (SCB), which consists of summary information intended to optimize the Files-11 space allocation, and the bit map itself, which lists the availability of individual blocks.

The SCB contains summary information about the volume (cluster factor, volume size, blocking factor, and so forth). Each bit in the bitmap represents an allocatable cluster on the volume. If a bit is set, the corresponding cluster is available for use. If a bit is clear, the cluster is not available.

During normal operation, the operating system moves portions of the bitmap in and out of cache memory. The state of each bit in memory is altered as clusters are allocated and deallocated. BITMAP.SYS is updated when the portion of the bitmap in cache is swapped back to disk. Since a portion of the bitmap is always in cache, BITMAP.SYS never reflects the current state of allocated clusters on a disk (unless the disk is dismounted or write-locked).

One of the functions of ANALYZE/DISK_STRUCTURE is to build a current version of BITMAP.SYS from data extracted from INDEXF.SYS, so that BITMAP.SYS accurately reflects the status of free clusters on the disk.

A.3.3 Bad Block File, BADBLK.SYS

The bad block file contains all the bad blocks on the volume. The system detects bad disk blocks dynamically and prevents their reuse once the files to which they are allocated have been deleted.

A.3.4 Master File Directory

The MFD is listed in the master file directory as 000000.DIR;1. The MFD, which is the root of the volume's directory structure, lists the reserved files that control the volume structure and may list both users' files and users' file directories.

Usually, however, the MFD is used to list the reserved files and users' file directories; users seldom enter files into the MFD, even on private volumes. In fact, on a private volume, it is most convenient for users to create a directory that has the same name as their default directory on a system disk. For an explanation of users' file directories and file specifications, see the *OpenVMS User's Manual*.

When the Backup utility (BACKUP) creates sequential disk save sets, it stores the save-set file in the MFD.

ANALYZE/DISK_STRUCTURE verifies all files contained in the directory structure by making comparisons to INDEXF.SYS. Any file found in INDEXF.SYS that is not traceable through the directory structure is "lost." ANALYZE/DISK_STRUCTURE places lost files in the top-level directory SYSLOST.DIR if you specified /REPAIR in the command.

A.3.5 Core Image File, CORIMG.SYS

The core image file is not used by the operating system.

A.3.6 Volume Set List File, VOLSET.SYS

The volume set list file is used only on relative volume 1 of a volume set. The file contains a list of the labels of all the volumes in the set and the name of the volume set.

ANALYZE/DISK_STRUCTURE uses VOLSET.SYS to locate each volume in the set and confirm the attributes of each volume. Since all volume set information is stored in VOLSET.SYS on relative volume 1, ANALYZE/DISK_STRUCTURE ignores VOLSET.SYS on all other volumes.

A.3.7 Continuation File, CONTIN.SYS

The continuation file is used as the extension file identifier when a file crosses from one volume to another volume of a loosely coupled volume set. This file is used for all but the first volume of a sequential disk save set.

A.3.8 Backup Log File, BACKUP.SYS

The backup log file is reserved for future use.

A.3.9 Pending Bad Block Log File, BADLOG.SYS

The pending bad block log file contains a list of suspected bad blocks on the volume that are not listed in the bad block file.

A.3.10 Quota File, QUOTA.SYS

The quota file is a reserved file that is used by the file system to keep track of the disk usage of each UIC on a volume. If you enable disk quota checking for a volume, the records of the file QUOTA.SYS contain all the UICs on the volume. The system constantly updates QUOTA.SYS to reflect the current disk usage, the maximum allowed disk usage, and the permitted overdraft for each UIC.

During the course of its operations, ANALYZE/DISK_STRUCTURE creates a version of QUOTA.SYS in memory that reflects the actual disk usage for each UIC. This version is eventually compared to the disk version of QUOTA.SYS. If ANALYZE/DISK_STRUCTURE detects any disparities in disk usage, ANALYZE/DISK_STRUCTURE notifies you. If you invoked ANALYZE/DISK_STRUCTURE with the /REPAIR qualifier, the disk version of QUOTA.SYS is updated.

A.3.11 Volume Security Profile, SECURITY.SYS

The volume security profile includes the volume owner UIC, the volume system-owner-group-world (SOGW) protection mask, and the volume access control list (ACL).

Files-11 Disk Structure

A.4 Files-11 ODS Level 1 Versus Level 2 (VAX Only)

A.4 Files-11 ODS Level 1 Versus Level 2 (VAX Only)

VAX

On VAX systems, for reasons of performance, reliability, and security, Files-11 ODS Level 2, a compatible superset of ODS Level 1, is the preferred disk structure on the system. At volume initialization time, Structure Level 2 is the default. (See the INITIALIZE command in the *OpenVMS DCL Dictionary*.)

On VAX systems, specify ODS Level 1 only for volumes that must be transportable to RSX-11M, RSX-11D, RSX-11M-PLUS, and IAS systems, as these systems support only that structure level. Additionally, you might be required to handle Structure Level 1 volumes transported to OpenVMS from one of these systems.

Where Structure Level 1 volumes are in use on the system, bear in mind the limitations on them that are shown in Table A-4.

Table A-4 Limitations on Files-11 Structure Level 1 Volumes

Disk	Only Files-11 ODS-2 disks are protected objects.
Directories	No hierarchies of directories and subdirectories, and no ordering of directory entries (that is, the file names) in any way. RSX-11M, RSX-11D, RSX-11M-PLUS, and IAS systems do not support subdirectories and alphabetical directory entries.
Disk quotas	Not supported.
Multivolume files and volume sets	Not supported.
Placement control	Not supported.
Caches	No caching of file header blocks, file identification slots, or extent entries.
System disk	Cannot be a Structure Level 1 volume.
VMScluster access	Local access only; cannot be shared across a cluster.
Clustered allocation	Not supported.
Backup home block	Not supported.
Protection code E	E means "extend" for the RSX-11M operating system but is ignored by OpenVMS.
File versions	Limited to 32,767; version limits are not supported.
Enhanced protection features (for example, access control lists)	Not supported.
Long file names	Not supported.
RMS journaling	Not supported.
RMS execution statistics monitoring	Not supported.

Future enhancements to OpenVMS software will be based primarily on Structure Level 2; therefore, Structure Level 1 volumes might be further restricted in the future. ♦

B

Tables of Time Differential Factors (TDFs)

The tables in this appendix show the TDFs of various locations in the world. Each table contains a list of locations in a specific region. The information in the tables is believed to be accurate at the time of publication.

Note

For the TDFs of countries in Africa, refer to the map in Figure 5-1.

Table B-1 lists the time differential factors for Europe.

Table B-1 TDFs for Europe

Region	Standard Time TDF	Daylight Saving Time TDF
Great Britain, Ireland	0:00	+1:00
Western European Time	0:00	+1:00
Iceland	0:00	—
Middle European Time	+1:00	+2:00
Poland	+1:00	+2:00
Eastern European Time	+2:00	+3:00
Turkey	+3:00	+4:00

Table B-2 lists the time differential factors for North America.

Table B-2 TDFs for North America

Region	Standard Time TDF	Daylight Saving Time TDF
U.S./Eastern Time	-5:00	-4:00
U.S./Central Time	-6:00	-5:00
U.S./Mountain Time	-7:00	-6:00
U.S./Pacific Time	-8:00	-7:00
U.S./Indiana (East)	-5:00	—
U.S./Alaska	-9:00	-8:00
U.S./Arizona	-7:00	—
U.S./Navajo	-7:00	-6:00

(continued on next page)

Tables of Time Differential Factors (TDFs)

Table B-2 (Cont.) TDFs for North America

Region	Standard Time TDF	Daylight Saving Time TDF
U.S./Michigan	-5:00	-4:00
U.S./Aleutian Islands	-10:00	-9:00
U.S./Hawaii	-10:00	—
U.S./Samoa	-11:00	—
Canada/Newfoundland	-3:30	-2:30
Canada/Atlantic	-4:00	-3:00
Canada/Eastern	-5:00	-4:00
Canada/Central	-6:00	-5:00
Canada/East-Saskatchewan	-6:00	—
Canada/Mountain	-7:00	-6:00
Canada/Pacific	-8:00	-7:00
Canada/Yukon	-9:00	-8:00

Table B-3 lists the time differential factors for Central and South America.

Table B-3 TDFs for Central and South America

Region	Standard Time TDF	Daylight Saving Time TDF
Mexico/BajaNorte	-8:00	-7:00
Mexico/BajaSur	-7:00	—
Mexico/General	-6:00	—
Cuba	-5:00	-4:00
Jamaica	-5:00	-4:00
Brazil/East	-3:00	-2:00
Brazil/West	-4:00	-3:00
Brazil/Acre	-5:00	-4:00
Brazil/DeNoronha	-2:00	-1:00
Chile/Regional	-4:00	-3:00
Chile/Easter Island	-6:00	-5:00

Table B-4 lists the time differential factors for Asia.

Table B-4 TDFs for Asia

Region	Standard Time TDF	Daylight Saving Time TDF
PRC (Mainland China)	+8:00	+9:00
ROK (Korea)	+9:00	+10:00

(continued on next page)

Tables of Time Differential Factors (TDFs)

Table B-4 (Cont.) TDFs for Asia

Region	Standard Time TDF	Daylight Saving Time TDF
Israel	+3:00	+4:00
Iran	+3:30	+4:30
Japan	+9:00	—
Singapore	+8:00	—
Hong Kong	+8:00	—
ROC (Taiwan)	+8:00	—

Table B-5 lists the time differential factors for the South Pacific.

Table B-5 TDFs for the South Pacific

Region	Standard Time TDF	Daylight Saving Time TDF
Australia/Tasmania	+10:00	+11:00
Australia/Queensland (standard time only)	+10:00	—
Australia/Queensland	+10:00	+11:00
Australia/North	+9:30	—
Australia/West	+8:00	—
Australia/South	+9:30	+10:30
Australia/Victoria	+10:00	+11:00
Australia/New South Wales	+10:00	+11:00
New Zealand	+12:00	+13:00

Table B-6 lists the time differential factors for Antarctica.

Table B-6 TDFs for Antarctica

Region	Standard Time TDF	Daylight Saving Time TDF
Antarctica	+0:00	—

Table 1: Descriptive Statistics

Variable	Mean (SD)	Range
Age	34.5 (10.2)	18-65
Gender		
Male	52.3%	
Female	47.7%	
Marital Status		
Married	68.9%	
Single	31.1%	
Relationship Duration	5.2 (3.8)	1-15
Frequency of Contact	3.1 (1.5)	1-5
Level of Contact	2.8 (1.2)	1-4

Note. SD = standard deviation.

Table 2: Correlation Matrix

Variable	1	2	3	4	5	6	7	8
1. Age	1.00							
2. Gender	-.05	1.00						
3. Marital Status	.12	.08	1.00					
4. Relationship Duration	.25	.15	.10	1.00				
5. Frequency of Contact	.30	.20	.18	.22	1.00			
6. Level of Contact	.28	.18	.15	.20	.25	1.00		
7. Perceived Support	.15	.10	.05	.12	.18	.15	1.00	
8. Psychological Distress	.10	.05	.02	.08	.12	.10	.15	1.00

Note. All correlations are significant at $p < .05$.

Table 3: Descriptive Statistics

Variable	Mean (SD)	Range
Age	34.5 (10.2)	18-65
Gender		
Male	52.3%	
Female	47.7%	
Marital Status		
Married	68.9%	
Single	31.1%	
Relationship Duration	5.2 (3.8)	1-15
Frequency of Contact	3.1 (1.5)	1-5
Level of Contact	2.8 (1.2)	1-4

Glossary

Following is an alphabetical listing of terms used in the *OpenVMS System Manager's Manual* and their definitions.

access control list (ACL)

A protection mechanism using a more refined level of protection than that available with **UIC-based protection**. ACLs can be used to grant or deny access to individual users or groups of users.

access mode

Any of the four processor access modes in which software executes. Processor access modes prevent system software from inadvertently performing operations that might damage the system. Processor access modes are in order from most to least privileged and protected: kernel, executive, supervisor, and user. When the processor is in any mode other than kernel mode, the processor is inhibited from executing privileged instructions.

account

Each system user has an account. When you log in, you log in under a particular account name and number. This number informs the system where your files are and what kind of access to other files and system facilities you should be given.

accounting files

Files where the system stores information on resource use. Compare with **current accounting file**.

active set

In a multiprocessing system, the subset of processors that have successfully run power-on diagnostics and are actively participating in system operations. Compare with **available set**.

active values

With system parameters, the set of values that is stored in memory and is used by the active system. When the system boots, it reads into memory the **current values** stored in a parameter file on disk.

adjacent node

In a network, a node that is connected to your node by a single physical line.

allocation class

In a VMScluster environment, for devices that are dual-ported between two computers, a numeric value used to create a unique, path-independent device name.

answer file

A file in the form SYS\$UPDATE:*product*.ANS. The file is created when you install a product initially, and you specify the Auto-Answer option. The file contains a record of the answers you entered when you ran VMSINSTAL.COM to install that product initially.

application service

A **LAT service** in which LAN users can access only a specific program. Contrast with **general timesharing service**.

area router

In a network, a node that performs routing operations between areas and within its own area. Also called a **level 2 router**. Compare with **level 1 router**.

autostart feature

A feature that simplifies startup and ensures high availability of execution queues in a VMScluster environment. It lets you do the following:

- Start all **autostart queues** on a node with a single command
- Specify a list of nodes (within a VMScluster environment) to which a queue can automatically fail over if necessary.

autostart queue

An execution queue that takes advantage of the **autostart feature**. When you create a queue, you can designate it as an autostart queue.

available set

In a multiprocessing system, those processors that have successfully completed the system's power-on hardware diagnostics and may or may not be actively involved in the system. Compare with **active set**.

backlink

In Files-11 disk structure, a pointer to the directory in which a file resides.

banner page

A specially formatted page that prints at the beginning and end of print jobs and files within print jobs. These pages are helpful in identifying and separating output jobs, and the files within those jobs, when they are printed.

base process priority

A base priority value that the system uses to schedule a process. Priorities range from a low of 0 to a high of 31; 0 through 15 are timesharing priorities and 16 through 31 are real-time priorities. Compare with **job scheduling priority**.

batch execution queue

An execution queue that can accept only batch jobs.

batch job

A detached process that sequentially runs one or more command procedures. The user defines the list of command procedures when submitting the job to a batch queue.

batch mode

An execution mode in which you can execute a command procedure by submitting the procedure to a batch queue. When resources are available, the system creates a detached process to execute the commands in the procedure. Usually, processes running in batch mode execute at a lower process priority, to avoid competing with interactive users for system resources.

beginning-of-tape (BOT) marker

A piece of photoreflexive tape that delimits the beginning of the writable area on a tape volume.

binding

On an InfoServer system, a function that creates a **virtual device unit** on a local OpenVMS system.

block

On Files-11 disks, the basic unit by which disk space is allocated (512 8-bit bytes). On magnetic tape, the size of a block is determined by the user.

boot block

Block 0 on a disk. It contains the location and size of the **primary bootstrap image**, which is used to boot the system. Certain processors, in order to boot, must read the boot block to obtain the location of the primary bootstrap image.

booting

Also called **bootstrapping**, the process of loading system software from the system disk into processor memory. You must install the operating system before you boot the system for the first time. See also **conversational boot** and **nonstop boot**.

bootstrapping

See **booting**.

bpi

Bits per inch; a measure used for characters of data on tape. Also called **density**.

caching

A performance enhancement in which the system stores information in memory; this includes information about a disk volume's free space, file identifications, quota file entries, and file headers.

capability

On VAX systems, software that makes the services of the vector processor available to system users.

circuit

In a network, a communications data path that connects **adjacent nodes**. A circuit is not a physical data path but, rather, a logical connection that operates over a physical connection (a line). All input and output (I/O) between nodes takes place over circuits.

cathedral windows

Segmented windows created from mapping windows; useful for reducing the overhead required to read large files. The Buffered I/O Byte Count Limit (BITLM) limits the number of cathedral windows a user can create.

cluster

On Files-11 media, a logical grouping of blocks; the basic unit by which disk space is allocated.

See also **VAXcluster system**, **VMScluster system**.

command procedure

A file containing DCL commands and, optionally, data used by those commands. When you execute a command procedure, the system reads the file and executes the commands it contains. This eliminates the need for you to enter each command separately. You can use command procedures to efficiently perform routine tasks. A command procedure can also be executed in **batch mode**.

command string

The complete specification of a command, including the command name, command qualifiers, parameters, and parameter qualifiers. Because a command can be continued on more than one line, the term is used to define the entire command.

Compact Disc Read-Only Memory (CD-ROM)

Computer discs similar to the CD-ROMs used for audio applications. The major difference is that CD-ROM computer disc players have a digital (rather than an audio) interface.

configuration database

In a network, each node has a configuration database that includes information about the node and other nodes with which it can communicate. The configuration database is made up of a **permanent database** and **volatile database**.

connection manager

In a VMScluster environment, the component that dynamically defines the VMScluster system and coordinates participation of computers in the cluster.

conversational boot

A booting operation in which you stop to perform special operations—for example, to change system parameter values—before booting. Contrast with **nonstop boot**.

Conversational boot operations are common in programming research and development environments where you must alter operating conditions for experimentation, testing, and debugging.

crash dump

When the operating system detects an unrecoverable error or an inconsistency within itself that causes the system to fail, it writes the contents of the error log buffers, processor registers, and memory into the **system dump file**.

crash history file

A file storing information about system crashes. Use the Crash Log Utility Extractor (CLUE) to display the contents of the crash history file to understand and resolve the issues responsible for crashes, and to obtain other useful data.

current accounting file

In a VMScluster environment, an **accounting file** for a particular node. By default, the current accounting file is SYS\$MANAGER:ACCOUNTNG.DAT.

current values

With system parameters, the set of values that is stored in the default parameter file on disk and are used to boot the system. When the system boots, it reads the current parameter values into memory to create **active values**.

cylinder

On a disk, consists of all **tracks** at the same radius on all recording surfaces of the disk.

data area

One of two divisions of CD-ROM volume space; includes the remaining volume space, beginning with logical sector 16.

DECevent

On Alpha systems, the event management utility that provides an interface between a system user and the operating system's event log files.

DECnet for OpenVMS

The name for the software and hardware products that allow various Digital operating systems to participate in a network. DECnet for OpenVMS allows a system to function as a node in a network.

default values

With system parameters, the set of values provided on your distribution kit and stored in the default list. These values allow you to boot any supported configuration.

density

A measurement, in bits per inch, used for characters of data on tape.

device

Hardware that allows access to storage media; also called **drive**.

device control library

A text library that contains user-written modules consisting of text or escape sequences. See also **device control module**.

device control module

A user-written module in a **device control library**. Device control modules can be used for the following purposes:

- With programmable printers, to insert device-dependent escape sequences that set up a printer for selected print options such as point size, character set, and bold or italic print.

- With both programmable and non programmable printers, to insert text at specific points in the processing of a print job.

See also **page setup module**, **reset module**, and **setup module**.

device driver

A system component that controls I/O operations for a particular device type. For a device to function on a system, the device must be connected and the device driver must be loaded into memory.

disk

Physical media on which files reside.

disk quota

A method for maintaining and enforcing limits on the amount of disk space available to users on a public volume. See also **quota file**.

drive

Hardware that allows access to storage media; also called **device**.

end node

In a network, a node that does not perform routing operations.

end-of-tape (EOT) marker

A piece of photorefective tape that delimits the end of the writable area on a tape volume.

ERRFMT process

System process that periodically empties the error log buffers, transforms the descriptions of the errors into standard formats, and stores the formatted information in the **error log file** on the system disk.

error log file

The operating system automatically records device and CPU error messages in this file. The Error Log utility invokes the **Error Log Report Formatter (ERF)** to selectively report the contents of an error log file.

Error Log Report Formatter (ERF)

A system component invoked by the Error Log utility to selectively report the contents of the **error log file**.

Ethernet

A single shared network channel, with all nodes having equal access to the channel. Ethernet offers local and remote connections as one integral network.

event classes

Categories of security-relevant events. The system always audits several event classes.

executable image

An image that can be run in a process. It is linked with the **/EXECUTABLE** qualifier (or without the **/SHAREABLE** qualifier) of the Linker utility.

execution queue

A queue that accepts batch or print jobs for processing. Compare with **generic queue**.

executive

A set of programs in the operating system that control the running of routines that perform I/O, resource allocation, and program execution. See also **executive routines**.

executive mode

The second most privileged processor **access mode**. OpenVMS Record Management Services (RMS) and many system service procedures execute in executive mode.

executive routines

System routines that detect errors and events and write relevant information into error log buffers in memory. See also **executive**.

expiration date

The Files-11 On-Disk Structure uses the expiration date of a file to track the use of a file. The expiration date aids in the disposal of seldom-used files.

extent

On Files-11 volumes, contiguous blocks allocated to a particular file.

feedback

Information, continuously collected by the **executive**, about the amount of various resources the system uses to process its work load. When run in feedback mode, AUTOGEN analyzes this information and adjusts the values for any related system parameters.

field

In a UAF record, a portion of the record you modify with the Authorize utility. The values you assign to each field do the following:

- Identify the user
- Define the user's work environment
- Control use of system resources

file

On Files-11 media, an array of consecutive virtual blocks, numbered 1 to n , plus a set of attributes with values. A file is either a data file or a directory file. Directories can contain both data files and directory files.

file banner page

A **banner page** that separates files within a job; users can override the file banner page settings you set for a queue.

file header

On a Files-11 volume, describes a portion of a file on the volume. File headers contain information such as the **owner UIC**, **protection code**, creation date and time, and **access control list (ACL)**.

file operation

In the Backup utility, an operation that processes individual files or directories.

Files-11 On-Disk Structure

A logical structure given to information stored on a disk; it is a hierarchical organization of files, their data, and the directories needed to gain access to them.

Files-11 volume

A disk volume that uses Files-11 On-Disk Structure and is mounted on a device.

full backup

See **image backup**.

full names

On VAX systems, hierarchically structured DECnet/OSI node names that can be stored in a DECdns naming service. Full names on VAX systems can be a maximum of 255 bytes long.

general timesharing service

A LAT service offering processing resources to users in the LAN. Contrast with **application service**.

generic batch queue

A generic queue that can direct jobs only to batch execution queues.

Generic batch queues are typically used in VMScluster environments to distribute the batch work load across several nodes.

generic output queue

A generic queue can direct jobs to any output execution queue. Generic output queues are typically used to distribute the output work load among several identical printers.

generic queue

A queue that holds batch or print jobs until they are transferred to an **execution queue** for processing.

A generic queue holds a job until an appropriate execution queue becomes available to initiate the job. The **queue manager** then requeues the job to the available execution queue.

group volume

A volume available to all the users in a group. Compare to **system volume**.

header labels

On magnetic tape, labels containing information such as the file name, creation date, and expiration date. When you create a file on magnetic tape, the magnetic tape file system writes header labels immediately preceding the data block. To access a file on magnetic tape by the file name, the file system searches the tape for the header label set that contains the specified file name.

header resident image

A **known image** for which the header of the image file remains permanently resident in memory, saving one disk I/O operation per file access.

home block

A block in a Files-11 volume that identifies it as a Files-11 volume. Usually, the home block is the next block after the **boot block** (block 0). If for some reason the home block cannot be read (is physically unusable), an alternative block is selected for use as the home block. This block provides specific information about the volume and default values for files on the volume.

identification record

A record of a **file header** that contains a summary of disk and volume characteristics.

image

A collection of procedures and data bound together by the Linker utility to form an executable program. Executable programs can be executed (or run) by a process. Usually, executable programs have the file type .EXE.

image backup

Also called a full backup. A Backup utility operation that saves a copy of all the files on a disk (or volume) to a special file called a **save set**. See also **image operation**.

image compare

A Backup utility operation that compares the contents of entire volumes.

image copy

A Backup utility operation that creates a new Files-11 On-Disk Structure on the output disk and copies an entire volume; the image backup is a logical duplicate of the contents of the disk.

image operation

A Backup utility operation that processes all files on the input disk.

image registry

A file associated with the Image Registry facility. To continue using a compatible application image that depends on a previous operating system version, you can register the image in the Image Registry.

image restore

A Backup utility operation that initializes the output disk and restores an entire volume.

incremental backup

A Backup utility operation that saves only those files that have been created or modified since the most recent backup that was performed using the /RECORD qualifier. (The /RECORD qualifier records the date and time that the files are backed up.)

incremental restore

A Backup utility operation that restores an incremental **save set**.

InfoServer system

An Ethernet-based, high-performance, **virtual device server**. The InfoServer system can serve physical device media and sets of logical disk blocks to client systems in a local area network (LAN). Systems running the appropriate client software can connect to virtual devices served by the InfoServer system and use them as though they are locally attached devices.

initialization file

In certain utilities, a file used each time you invoke the utility. In the initialization file, you can perform tasks such as defining keys and setting up your environment.

installation procedure

The procedure for installing the operating system for the first time. Also, a procedure for installing a layered product.

IRG (interrecord gap)

On magnetic tape, the interval of space between blocks.

job banner pages

banner pages that identify jobs; users cannot override job banner pages that you set for a queue. Compare with **file banner pages**.

job controller

The system process that creates a process to perform the tasks in a batch job.

job scheduling priority

A priority value that the system uses to schedule a batch or print jobs in a queue. Job scheduling priorities range from a low of 0 to a high of 255. Compare with **base process priority**.

kernel mode

The most privileged processor **access mode**. The operating system's most privileged services, such as I/O drivers and the pager, run in kernel mode. When in kernel mode, the processor has complete control of, and responsibility for, the system.

known file list

An internal data structure on which the system defines **known images**. Each entry in the known file list identifies the file name of the known image and the attributes with which it was installed.

known image

An image installed with the Install utility (INSTALL). When you install an image, the image is assigned attributes and becomes known to the system.

LASTport protocol

A specialized LAN transport protocol, implemented by the InfoServer software, that allows many clients to access InfoServer systems and perform reliable device read and write operations.

The LASTport/DISK protocol and LASTport/TAPE protocol are specialized disk and tape protocols that use the LASTport protocol.

See also **InfoServer system**.

LAT protocol

Protocol, implemented by the LAT software, that allows the operating system to offer resources, or LAT services that **terminal servers** can access.

LAT service announcements

Multicast messages sent by **LAT service nodes** and used to create a database of service nodes available.

LAT service node

A system that supports incoming LAT connections or a system that offers **LAT services**.

LAT services

Computing resources made available to users in the LAN through the LAT software. A LAT service can be a **general timesharing service** or an **application service**.

level 1 router

In a network, a node that performs routing operations within a single area. Compare with **level 2 router**.

level 2 router

In a network, a node that performs routing operations between areas and within its own area. Also called an **area router**. Compare with **level 1 router**.

license

Many software vendors provide software to their customers under an agreement called a license. Although the term *license* can have specific legal connotations, for the purpose of this manual a license refers to the authorization you have to use a product.

The License Management facility (LMF) lets you register, manage, and track software licenses on line. See also **Product Authorization Key (PAK)**.

line

In a network, a physical data path that connects **adjacent nodes**. A communications line connects your computer to the DECnet for OpenVMS network.

load address

The location in memory (specified in hexadecimal notation) to which the system loads the **bootstrap image**.

Local Area VAXcluster configuration

A VAXcluster configuration in which a single VAX computer serves as the management center of the cluster, plus one or more VAX computers that are connected to this hub.

local cluster

In the System Management utility (SYSMAN), the node from which you are executing SYSMAN.

local node

In a network, the node on which you are working.

In the System Management utility (SYSMAN), the node on which you execute SYSMAN.

Contrast with **remote node**.

logical block

Organizational unit of volume space. The logical block size cannot exceed the logical sector size.

logical block numbering

Begins with the first byte in the volume space and continues in a sequentially ascending order through the remainder of the volume space.

logical link

In a network, connects two processes and carries a stream of two-way communications traffic between the processes over a **circuit**. A single circuit established between two nodes can support many logical links concurrently.

logical queue

A special type of generic output queue that transfers print jobs to another output execution queue. You might use this kind of queue to temporarily redirect a queue when the device on which it runs is broken.

logical sector

Organizational unit of a volume; consists of one or more physical sectors. No more than one logical sector can begin in any physical sector.

Logical sectors are numbered in ascending order, with 0 assigned to the logical sector having the lowest physical address containing recorded data. Each logical sector includes a data field made up of 2048 or more bytes (the number of bytes always equals a power of 2).

login command procedure

A command procedure that executes each time a user logs in. Add commands to a login command procedure to execute commands when a user logs in, for example, to set up the user environment.

login (LGI) system parameters

System parameters that control login functions. The names of these system parameters begin with LGI.

loopback tests

In a network, a series of tests to help determine whether the network is operating properly.

lost file

A file that is not linked to a directory. When you delete a directory file (a file with the file type .DIR) without first deleting its subordinate files, the files referred to by that directory become lost files. Lost files are a nonproductive use of disk space and act as debits against a user's **disk quota**.

Magnetic Tape Ancillary Control Process (MTACP)

The internal software process of the operating system that interprets the logical format of standard labeled tape volumes.

maintenance release

A release of the operating system that is applied with an **update procedure**.

mandatory update

A software update that is required immediately after upgrading or installing the operating system.

mass storage control protocol (MSCP) server

In a VMScluster environment, the component that implements the MSCP protocol, which is used to communicate with a controller for DSA disks, such as RA-series disks. In conjunction with one or both of the disk class **device drivers** (DUDRIVER, DSDRIVER), the MSCP server implements this protocol on a computer, allowing the computer to function as a storage controller.

master file directory (MFD)

The file that contains the name of all user file directories on a disk.

media

The physical substance on which you can store data.

mount verification

A recovery mechanism for disk and tape operations. If a device goes off line or is **write-locked** while mount verification is enabled, you can correct the problem and continue the operation.

multivolume file

A file that is continued on another volume when the data blocks of a file or related files do not physically fit on one volume (a reel of magnetic tape).

network

A means of connecting computers that allows them to share or transfer information or communications. A network includes two or more computers that are connected, and the hardware and software that makes those connections.

network proxy account

A user account that allows users on a **remote node** in a network to access data by way of a local account on your system. Proxy accounts are useful when you want to grant one or more users on a remote node access to specific files but you do not want to give them a private account on your system.

node

In a network, a computer system that is connected to another system in a network—by means of cables, telephone lines, microwave and satellite links, for example.

nonlocal cluster

In the System Management utility (SYSMAN), any cluster other than the one from which you are executing SYSMAN.

nonlocal environment

In the System Management utility (SYSMAN), your environment when you are not working on your local node or within your own cluster.

nonstop boot

The most common booting operation. You perform a nonstop boot if you do not want to stop to perform special operations—for example, to change system parameter values—before booting. Contrast with **conversational boot**.

object

In a network, a process to which a **logical link** connects. Some objects are DECnet programs—for example, the MAIL object; other objects are user-written programs.

For two programs to communicate over the network, the source program on the local node establishes a **logical link** with the object on the remote node.

OPCOM messages

Messages broadcast by the Operator Communication Manager (OPCOM). These messages are displayed on **operator terminals** and written to the **operator log file**. The messages might be general messages that you send, user requests, operator replies, or system events.

OPCOM process

The system process that manages Operator Communication Manager (OPCOM) operations.

operator log file

The Operator Communication Manager (OPCOM) records messages in this file. The file is named SYS\$MANAGER:OPERATOR.LOG.

operator terminals

Terminals designated to display messages broadcast by the Operator Communication Manager (OPCOM). Usually, the console terminal (with the device name OPA0:) is the operator terminal. However, you can designate any user terminal as an operator terminal.

output execution queue

A queue that accepts jobs for processing by a **symbiont**. The **queue manager** sends the symbiont a list of files, which the user defines when submitting the job. An output symbiont transfers data from a disk to an output device. As the symbiont processes each file, it produces output for the device it controls, such as a printer or a terminal.

owner UIC

Used with **UIC-based protection**, usually the UIC of the person who created a file or volume.

page

A unit used for allocating and deallocating memory.

On VAX systems, a page is 512 bytes.

On Alpha systems, a page can be 8 kilobytes (KB) (8192 bytes), 16 KB, 32 KB, or 64 KB. The initial set of Alpha computers use a page size of 8192 bytes. Compare with **pagelet**.

page file

In a **paging** operation, the file to which the system writes paged portions of memory. Your distribution kit includes a page file named **SYS\$SYSTEM:PAGEFILE.SYS**. If necessary, **SYS\$SYSTEM:PAGEFILE.SYS** can be used in place of the system crash dump file.

pagelet

On Alpha systems, a unit of memory in a 512-byte quantity. One Alpha pagelet is the same size as one VAX page. Also, on an Alpha 8KB computer, 16 Alpha pagelets equal 1 Alpha page.

page setup module

A **device control module** inserted at the beginning of each page of a print job.

paging

A memory management operation to efficiently use the physical memory allotted to a process by moving information between physical memory and files stored on disk. In paging, the system moves infrequently used portions of a process workspace out of physical memory to a file. Compare with **swapping**.

PAK

See **Product Authorization Key (PAK)**.

partition

A logical subset of a read/write disk. A single disk can be subdivided into several partitions, each of which can be used independently. The partitions appear to be whole disks.

permanent database

In a network, a permanent copy of the DECnet for OpenVMS **configuration database**. When you start the network, the permanent database provides the initial values for the **volatile database**. Changes remain after the network is shut down, but do not affect the current system.

permanently open image

A **known image** where directory information on the image file remains permanently resident in memory, eliminating the usual directory search required to locate a file.

physical dump

A **crash dump** containing the entire contents of physical memory to the **system dump file**. Compare with **selective dump**.

physical operation

In the Backup utility, an operation that copies, saves, restores, or compares an entire volume by logical blocks, ignoring any file structure.

physical sector

Division of a system or data area; smallest addressable unit on an ISO 9660 CD-ROM.

primary bootstrap image

Program that the boot block points to, which allows access to the system disk by finding the the **secondary bootstrap image**, SYSBOOT.EXE, and loading it into memory.

On VAX systems, the primary bootstrap image is VMB.EXE.

On Alpha systems, the primary bootstrap image is APB.EXE.

primary page and swap files

The default **page file** and **swap file** provided with your distribution kit. These files are named SYS\$SYSTEM:PAGEFILE.SYS and SYS\$SYSTEM:SWAPFILE.SYS. Contrast with **secondary page and swap files**.

primary processor

In a multiprocessing system, the processor that is either logically or physically attached to the console device and is the target of the console commands that bootstrap the multiprocessing system. The primary processor is responsible for starting other processors in the multiprocessing system. It also serves as the system timekeeper.

print forms

You can use print forms with output queues to determine certain page formatting attributes (such as margins and page length). In addition, the paper stock specified in a form determines whether a job is printed; if the stock of a job's form does not match the stock of the form mounted on the queue, the job is not printed.

Digital supplies a default print form named DEFAULT. You can create additional forms if users need help formatting output, or if certain print jobs require special paper.

print job

An entry in an output queue that specifies a file or files to be printed on a printer. The user defines the file or files to be printed when submitting the job. When a printer is available, the **queue manager** sends the file to a **symbiont** for formatting and printing.

printer queue

A type of output execution queue that uses a **symbiont** to direct output to a printer. Compare with **server queue** and **terminal queue**.

priority

See **base process priority** or **job scheduling priority**.

private volume

A file-structured disk volume that contains only private files.

privileged image

A **known image** where increased **privileges** are temporarily assigned to any process running the image, permitting the process to exceed its user authorization file (UAF) privilege restrictions during execution of the image. In this way, users with normal privileges can run programs that require higher-than-normal privileges.

privileges

A means of restricting the functions users are authorized to perform on the system. System managers require privileges that are denied to most users.

process limits and quotas

User authorization file (UAF) parameters you can set for a user account to control the usage of system resources by processes in that account. (UAF parameters are different than system parameters.) You set values for process limits and quotas using the Authorize utility.

Product Authorization Key (PAK)

Information, typically on a piece of paper, provided for many Digital products. The data provided in the PAK allows you to register a software **license** in the license database on a system.

protected image

A **known image** that is a **shareable image** and contains protected code. Protected code is code that runs in **kernel mode** or **executive mode** but that can be called by a **user mode** image.

protection code

Used with **UIC-based protection**, indicates who is allowed access and for what purposes.

public volume

A Files-11 volume that any user on the system can access and that can contain both private and public files.

queue

Allows users to submit requests for printing or batch processing. The system prints users' print jobs or processes users' batch jobs as resources allow.

queue characteristics

Characteristics you can define and assign to a queue To control the batch or print jobs that execute on the queue.

queue database

A file or files that store information about queues and batch and print jobs.

queue manager

The system component that controls queue activity.

quota file

On Files-11 volumes, the file that records all users who are allowed to use a disk and that shows their current disk usage and their maximum disk allocation. A quota file, QUOTA.SYS, which is stored in directory [000000] with other system files, requires 1 block of disk storage for every 16 entries. See also **disk quotas**.

record blocking

On Files-11 volumes, the grouping of individual records into a block, thereby reducing wasted space.

remote node

In a network, a node that is accessible to the node you are working on (the local node) over the network.

In the System Management utility (SYSMAN), any node other than the one on which you are executing SYSMAN.

Contrast with **local node**.

reset module

A **device control module** inserted at the end of each print job. Use reset modules to reset a printer at the end of a job.

resident image

On Alpha systems, a **known image** that improves the performance of a **shareable image**. With a resident image, portions of images that contain code are moved into system space, where they reside on a large single page, thus improving performance.

root volume

The first volume in a **volume set**. Each volume in the volume set is identified by a volume number relative to the root volume, which is always relative to volume 1.

router

In a network, a node that performs routing operations.

routing

In a network of more than two nodes, the process of directing a data message from a source node to a destination node (known as an **end node**). Both routers and end nodes can send messages to and receive messages from other nodes in the network.

ruleset

Software routine or function that is analogous to an executable file; used by DECEvent.

save set

A special file used by the Backup utility. The Backup utility saves files to a save set and restores files from a save set. Installation and upgrade procedures restore product files from a save set to your system disk.

scalar

A single data item, having one value. Compare with **vector**.

secondary bootstrap image

Image that allows access to the system disk: SYS\$SYSTEM:SYSBOOT.EXE.

secondary page and swap files

Additional **page files** and **swap files** that you might create for performance or disk space reasons. The system uses the space in the secondary files for paging and swapping in addition to the space in the **primary page and swap files**.

secondary processor

In a multiprocessing system, any processor that is not a **primary processor**.

sector

The smallest unit discernible to the Files-11 On-Disk structure. For most Files-11 disks, a sector is equivalent to a block (512 bytes).

On ISO 9660 volumes, a uniquely addressable unit; each sector on a CD-ROM comprises a sequence of 2048 8-bit bytes.

security audit log file

A clusterwide file that contains a record of security events on the system. Using the ANALYZE/AUDIT command, you can produce reports and summaries of security events from the security audit log file.

selective dump

A **crash dump** containing only those portions of memory most likely to be useful in a crash dump analysis. A selective dump is useful when sufficient disk space is not available to hold all physical memory. Compare with **physical dump**.

selective operation

A Backup utility operation that processes files or volumes selectively, according to criteria such as version number, file type, UIC, date and time of creation, expiration date, or modification date.

sequential organization

On magnetic tape media, the organization of data; that is, data is organized in the order in which it is written to the tape.

server queue

A type of output execution queue that uses a user-modified or user-written **symbiont** to process the files that belong to print jobs in the queue. Compare with **printer queue** and **terminal queue**.

setup module

A **device control module** inserted at the beginning of a file in a print job.

system image snapshot

A record of the system setup used with the Snapshot facility.

system messages

Messages returned by the system when you enter commands in DCL or in utilities. These messages help you understand the result of each command.

system parameters

Parameters for which you can set values to control how the system functions. Values of system parameters control a wide range of system functions including but not limited to memory management, process scheduling, and system security.

system volume

A volume available to all the users on a system. Compare to **group volume**.

systemwide logical name

A logical name that applies to the entire system. It is defined in the system logical name table and can be used by any process in a system.

tape mass storage control protocol (TMSCP) server

In a VMScluster environment, the component that implements the TMSCP protocol, which is used to communicate with a controller for local MSCP tapes, such as TU-series tapes. In conjunction with the tape class **device driver** (TUDRIVER), the TMSCP server implements this protocol on a processor, allowing the processor to function as a storage controller.

target disk

In VMSINSTAL.COM or VMSKITBLD.COM, the disk to which you move the system files. Compare with **source disk**.

terminal queue

A type of output execution queue that uses a **sybiont** to direct output to a terminal printer. Compare with **printer queue** and **server queue**.

terminal servers

Communication devices dedicated for connecting terminals, modems, or printers to a local area network (LAN) and to other systems within a LAN. See also **LAT protocol**.

track

On a disk, the collection of **sectors** (or blocks, on Files-11 volumes) at a single radius on one recording surface of the disk. It is accessible to a given read/write head position on the disk device.

trailer labels

On magnetic tape, labels similar to **header labels**, but written following the file.

trusted logical names

Logical names associated with **executive mode** or **kernel mode**.

tuning

The process of altering various system values to obtain the optimum *overall* performance possible from any given configuration and work load.

UAF

See **user authorization file (UAF)**.

UETP (User Environment Test Package)

A software package designed to test whether the OpenVMS operating system is installed correctly.

UIC

See **user identification code (UIC)**.

UIC-based protection

A protection mechanism based on the **user identification code (UIC)** and applied to all protected objects. Compare with **access control list (ACL)**.

update procedure

Procedure used if you have a previous version of the operating system and you want to make minor fixes to it. When you update the operating system, the update procedure replaces some system files.

upgrade procedure

If you are already running a standard version of the operating system, you can use the upgrade procedure to obtain a higher version.

user authorization file (UAF)

A file containing an entry for every user that you authorize to gain access to the system. Each entry identifies the user name, password, default account, UIC (user identification code), quotas, limits, and privileges assigned to individuals who use the system.

User Environment Test Package (UETP)

See **UETP**.

user identification code (UIC)

The pair of numbers assigned to users, files, and other system objects, that specify the type of access available to the owner, group, world and system. The UIC consists of a group number and a member number separated by a comma and enclosed within square brackets. Same as UIC. See also **account** and **UIC-based protection**.

user mode

The least privileged processor **access mode**. User processes and run-time library routines run in user mode.

utility program

A program supplied by Digital that performs a set of related operations. For example, the Backup utility (BACKUP) allows you to save and restore files.

- AUTOGEN feedback, 14-8 to 14-22
 - checks performed on, 14-11
 - collection of, 14-11
 - examining effect on parameters, 14-11
 - file stored in, 14-11
 - importance of system work load, 14-11
 - improving system performance, 14-10
 - maximum age, 14-11
 - minimum age, 14-11, 14-21
 - report file, 14-11
 - sample, 14-12
 - sending automatically, 14-8, 14-22
 - resources affected by, 14-10
 - saving during system shutdown, 4-29
- Automatic configurations
 - of DECnet node, 21-12
 - of devices, 7-6
- Automatic login facility (ALF)
 - setting up an automatic login account, 6-32
- Automatic start
 - See also Autostart feature
 - of queue manager, 12-3, 12-7, 12-9
- Automatic volume switching, 8-39
- Autostart feature
 - See also Autostart queues
 - description, 13-4
 - disabling, 12-9, 13-54
 - before shutting down a node, 13-55
 - enabling, 13-4, 13-18, 13-48
 - on LAT queues, 13-4, 13-10
 - recommended use, 13-10
 - with LAT queues, 13-4, 13-10
- Autostart queues
 - See also Autostart feature
 - activating, 13-4, 13-15, 13-16, 13-49
 - activating inactive queue status, 13-51
 - creating, 13-15, 13-16
 - preventing from starting, 13-54
 - recommended use, 13-10
 - relationship between activating and starting, 13-4
 - starting, 13-48
 - in startup command procedure, 13-18
 - troubleshooting, 13-81
 - with LAT printers, 13-4, 13-10
- AUTO_POSITIONING command
 - SHOW CLUSTER, 20-12
- Availability
 - of queue manager, 12-9, 12-17
 - of queues, 13-4, 13-15
- Availability of devices
 - OPCOM message, 8-59
- Available queue status, 13-51
- Available sets, 26-2

B

- Backing up the system disk, 5-49
 - after installation, 3-13
- Backlinks
 - definition, 8-56
- BACKUP.SYS file
 - See Backup log file
- BACKUP command
 - and save sets, 10-23
 - /EXACT_ORDER qualifier, 10-22
 - for backing up directories, 10-24
 - for copying directories, 10-23
 - for copying files, 10-23
 - for image backups, 10-31, 10-32
 - for incremental backups, 10-34, 10-35
 - format, 10-4
 - /GROUP_SIZE qualifier, 10-62
 - /IGNORE=LABEL_PROCESSING qualifier, 10-22, 10-65
 - /IGNORE qualifier, 10-30, 10-63
 - /IMAGE qualifier, 10-31, 10-32, 10-55
 - /INITIALIZE qualifier, 10-14
 - in VMSINSTAL.COM command procedure, 3-13
 - /JOURNAL qualifier, 10-26
 - /LABEL qualifier, 10-13, 10-23, 10-34
 - /LOG qualifier, 10-63
 - /PHYSICAL qualifier, 10-55
 - qualifiers, 10-5
 - /RECORD qualifier, 10-31, 10-32, 10-34, 10-35
 - /REWIND qualifier, 10-13, 10-34
 - /SAVE_SET qualifier, 10-32, 10-35
 - /SINCE qualifier, 10-34, 10-35
 - /VERIFY qualifier, 10-63
 - with multiple output devices, 10-24, 10-32, 10-33
- Backup log file
 - BACKUP.SYS, A-9
 - reserved file, A-9
- Backup Manager, 10-5
 - features, 10-5
 - getting started, 10-6
 - types of help available, 10-5
- BACKUP media
 - Files-11 disk save set, 10-7
 - magnetic tape save set, 10-6
 - network save set, 10-7
 - sequential-disk save set, 10-7
- Backups
 - image
 - See Image backup
 - incremental
 - See Incremental backup
 - performing in command procedures, 8-46

Backups (cont'd)

standalone

See Standalone BACKUP

Backup utility (BACKUP)

backing up active disks, 10-63

backing up shadow sets, 10-40

Backup Manager, 10-4

command format, 10-4

command procedures, 10-36

command qualifiers, 10-5

copying dump file, 15-4, 15-14

copying the queue database, 12-12

data integrity mechanisms, 10-62

InfoServer tapes, 10-60

interfaces to it, 10-4

open files during a backup, 10-30, 10-63

restoring shadow sets, 10-47

save set, 10-23

stores save-set file in MFD, A-8

transferring information, 9-19

use by VMSINSTAL.COM command procedure, 3-13

using on workstations, 10-36

using to back up directories, 10-24

using to copy directories, 10-23

using to copy files, 10-23

BACKUSER.COM command procedure, 10-36

BADBLK.SYS file

See Bad block file

Bad block file

BADBLK.SYS, A-8

definition, A-8

reserved file, A-8

Bad Block Locator utility (BAD)

analyzing media for bad blocks, 8-65

detecting media errors, 8-64

invoking with ANALYZE/MEDIA, 8-65

BADLOG.SYS file

See Pending bad block log file

Banner pages

commands used with, 13-60

definition, 13-33

file, 13-42, 13-60

job, 13-42, 13-60

BASEENVIRON phase of system startup, 5-4, 5-18

Base priority

See also Priority, base

Batch and print queuing system, 12-2

See also Queue configurations
components, 12-1

in VMScluster environments

with multiple system disk, 12-6

queue database

location of files, 12-5

mounting disk for, 12-6

queuing process, 13-2

Batch and print queuing system (cont'd)

sample configurations, 13-4 to 13-14

starting in system startup, 5-12

steps for setting up, 13-14

Batch execution queues

See also Execution queues

description, 13-3

Batch identifiers, 11-11

Batch jobs

See also Batch processing environment

See also Batch queues

accessing devices, 8-46

allowing to complete before stopping a queue, 13-55

changing scheduling priority, 13-72

completing before using VMSINSTAL.COM

command procedure, 3-7

controlling, 13-68

deleting, 13-74

distributing system work load, 16-4

executing, 13-3

holding and releasing, 13-71

job card, 7-17

modifying, 13-70

monitoring, 13-68

requeuing

executing, 13-73

pending, 13-73

retaining in a queue, 13-74

scheduling, 13-72

status

See Job status

submitting at startup, 5-14

Batch mode

as execution mode for a startup procedure, 5-19

Batch processing environment

See also Batch jobs

See also Batch queues

generic queues in a VMScluster, 13-7

on a standalone workstation, 13-5

sample configurations, 13-4 to 13-7

steps for setting up, 13-14

with specialized queues, 13-6

Batch queues

See also Batch jobs

See also Batch processing environment

allowing jobs to complete before stopping, 13-55

commands for managing, 13-46

creating, 13-15

deleting, 13-57

for memory constrained systems, 13-31

on standalone workstations, 13-5

optimizing for Sort/Merge utility, 13-31

options, 13-18

characteristics, 13-28

- Batch queues
 - options (cont'd)
 - controlling job performance and resources, 13-28
 - qualifiers for specifying, 13-19 to 13-20
 - restricting access, 13-22
 - retaining jobs, 13-26
 - pausing, 13-53
 - setting up for vector processing, 26-7
 - starting, 13-15
 - status, 13-50
 - stopping, 13-54, 13-55
 - before shutting down a node, 13-55
- Binding volumes into a volume set, 8-22
- BIOLM process limit, 6-41
 - value for efficient backups, 10-10
- BITMAP.SYS file
 - See Storage bit map file
- Blocks
 - disk
 - definition, 8-2
 - erasing, 8-58
 - Files-11
 - definition, A-1
 - tape
 - definition, 8-6
- Boot block
 - definition, 4-17
 - in index file, A-6
 - processors using, 4-17
 - role in boot process, 4-2
 - writing with Writeboot utility, 4-17
- Booting
 - automatically after shutting down, 4-29
 - bootstrap image
 - Alpha, 4-18
 - in index file, A-6
 - VAX, 4-18
 - conversationally
 - See Conversational boot
 - cross-architecture, 22-23
 - definition, 4-2
 - displaying startup commands, 4-14
 - from an alternate system disk, 4-3
 - in a multiprocessing system, 26-2
 - in an emergency
 - with default system parameters, 4-8
 - without startup and login procedures, 4-8
 - without the user authorization file, 4-10
 - installation of page and swap files, 15-5
 - loading VVIEF code, 26-5
 - location of computer-specific instructions, 4-2
 - message
 - See Boot messages
 - nonstop, 4-3
 - problems

- Booting
 - problems (cont'd)
 - fixing by booting with default parameter values, 4-8
 - fixing by booting with minimum startup, 4-14
 - hardware, 4-16
 - invalid boot block, 4-17
 - solving, 4-16
 - use of boot block, 4-17
 - with an alternate system parameter file, 4-7
 - with controlled system startup, 4-12
- Boot messages
 - indicating execution of STARTUP.COM
 - procedure, 4-5
 - indicating execution of STARTUP_VMS.COM
 - procedure, 4-5
 - indicating success, 4-5
 - indicating that login is possible, 4-5
 - question mark (?), 4-16
- Bootstrapping
 - See Booting
- BOT (beginning-of-tape)
 - See BOT markers
- BOT markers, 8-6
- Break-ins
 - auditing attempts, 18-25
- Bridges
 - network, 21-7
- BROADCAST device setting, 18-23
- Buffered I/O
 - byte count limit, 6-42
 - count limit, 6-41
- Bugcheck message
 - during UETP, 17-30
- Burst bars, 13-33
- Burst pages, 13-33
 - file, 13-35
 - job, 13-33
- Busy queue status, 13-51
- BYTLM process limit, 6-42
 - value for efficient backups/nomaster, 10-10

C

- C2-class system
 - installing OpenVMS to use, 3-2
- CACHE_SERVER process
 - creation during system startup, 5-5
- Caching
 - ACP system parameters, 8-43
- Canceling
 - characteristics on a queue, 13-59
- Capability of a vector
 - See Vector capability

- Captive accounts, 6-29
- Card readers
 - operating, 7-17
 - translation modes, 7-18
- Cards
 - decks, 7-17
- CD-ROM (compact disc read-only memory), 8-4, 8-5
 - accessing ISO 9660-formatted, 8-34
 - accessing with the InfoServer Client for OpenVMS, 23-10
 - characteristics of, 8-4
 - data arrangement on, 8-5
 - file structures, 8-4
 - formats for automatic serving, 23-2
 - High Sierra format, 8-4
 - ISO 9660 format, 8-4
 - media formats used, 8-4
- Changing Digital-supplied data in the Help Message database, 5-31
- Changing scheduling priority for a batch or print job, 13-72
- Changing size of page, swap, and dump files
 - recommended method, 15-22
- Changing system parameters
 - recommended method, 14-4, 14-18
 - with AUTOGEN, 14-18
 - with conversational boot, 4-6, 14-36
 - with SYSGEN, 14-33
 - with SYSMAN, 14-28
- Changing the DEFAULT form, 13-63
- Changing the system parameter file
 - with SYSGEN, 14-33
 - with SYSMAN, 14-28
- Channels
 - network, 21-5
- CI (computer interconnect)
 - data link between cluster nodes, 21-8
- Circuits
 - network, 21-5
 - definition, 21-4
- Classes
 - enabling and disabling, 18-23
- Class of data
 - in SHOW CLUSTER, 20-4
 - removing, 20-11
- CLEAR command
 - NCP, 21-11
- Closed queue status, 13-51
- Closing a queue, 13-53
- Clusters
 - See also VAXcluster environments
 - See also VMScluster systems
 - See VAXcluster environments
 - See VMScluster systems
- CLUSTER_CONFIG.COM command procedure, 5-6, 20-2
 - compared to VMSKITBLD.COM command procedure, 2-31
 - creating SATELLITE_PAGE.COM command procedure, 15-20
 - setting up LAN MOP, 22-21
 - use in adding a system to a VMScluster, 2-31
- CLUSTER_SERVER process
 - creation during system startup, 5-5
- CLUSTER_SIZE attribute, 10-57
- Command files
 - LAN Control Program (LANCP) utility, 22-9
- Command formats
 - for backups, 10-4
 - for image backups, 10-31, 10-32
 - for incremental backups, 10-34, 10-35
 - for multiple backup output devices, 10-24, 10-32, 10-33
- Command procedures
 - executing in SYSMAN, 2-18
 - executing with SYSMAN DO command, 2-18
 - for backups, 10-36
 - for image backups, 10-36
 - for incremental backups, 10-38
 - for installing products, 3-32
 - for interactive backups, 10-39
 - for MONITOR
 - archiving recording and summary files, 18-39
 - creating cluster summaries, 18-40
 - creating summary file, 18-39
 - generating clusterwide multfile summary reports, 18-39
 - initiating continuous recording, 18-41
 - invoking SUBMON.COM from startup, 18-39
 - MONITOR.COM, 18-39
 - MONSUM.COM, 18-39
 - rerecording monitoring, 18-39
 - starting MONITOR.COM as a detached process, 18-39
 - SUBMON.COM, 18-39
 - for POLYCENTER Software Installation utility, 3-18
 - for SHOW CLUSTER, 20-15
 - controlling output, 20-5
 - default file type, 20-16
 - description, 20-16
 - formatting reports, 20-10
 - initialization, 20-14
 - SHOW_CLUSTER\$INIT.COM, 20-14
 - for system management (overview), 2-5
 - for system startup, 2-11
 - See also Startup command procedure
 - login
 - setting protection in, 9-7
- LOGIN.COM, 2-16

Command procedures (cont'd)

- setting up storage media, 8-46
 - disk volumes, 8-46
 - tape volumes, 8-46
- testing a spooled printer, 7-16
- to configure a DECnet network node, 21-11
- to install software products
 - See VMSINSTAL.COM command procedure
- to run AUTOGEN regularly, 14-22
- to specify default state of operator log files, 18-23
- VVIEF\$INSTAL.COM procedure, 26-5

Commands

- See also DCL commands
- executing on multiple nodes, 2-18
- executing on remote nodes, 2-12
- executing with SYSMAN DO command, 2-18

Communicating

- with operators, 8-18
- with users, 8-18
 - creating systemwide announcements, 5-14

Communications line

- definition, 21-12
- network
 - definition, 21-4

Compact disc drives

- supported by UETP, 17-7

Compact disc read-only memory

- See CD-ROM

Compare operations

- with BACKUP, 10-25

Completion status

- showing for batch and print jobs, 13-26

Compressing the Help Message database after deletions, 5-30

Computer interconnect

- See CI

Configuration databases

- building manually using NCP, 21-11
- changing location of file, 21-10
- contents of, 21-10
- modifying using NCP, 21-11
- permanent database, 21-11
- stored in SYS\$SYSTEM:NETNODE_REMOTE.DAT file, 21-10
- volatile database, 21-11

Configurations

- displaying LAN, 22-9
- for software product options, 3-22
- LAN, 22-1
- LAN firmware updates, 22-14
- network
 - DECnet support, 21-6
 - local area LAN, 21-6
 - multiple-area, 21-6
 - single-area, 21-6
- queue

Configurations

queue (cont'd)

- sample batch queuing system, 13-4 to 13-7
- sample print queuing system, 13-8 to 13-14

CONFIGURATION SET command

- in SYSMAN, 20-16

CONFIGURATION SHOW command

- in SYSMAN, 20-16

CONFIGURE phase of system startup, 5-4, 5-18, 7-8

- definition, 7-6

CONFIGURE process

- starting, 5-5

Configuring DECnet, 21-7

- configuration database, 21-10
- nodes

- automatic, 21-12

- manual, 21-12

- planning network, 21-12

Configuring devices

HSC

- disabling during system startup, 7-8
- in system startup, 7-6
- in system startup, 5-7, 7-6

CONINTERR.EXE driver, 7-7

CONNECT command

- See also IO CONNECT command

- in SYSGEN (VAX), 7-7

- for connecting the console device, 7-7
- in system startup, 5-7

Connecting devices

- automatically, 7-6, 7-7

- in system startup, 5-4

- manually, 7-8

- in system startup, 5-7

- on VAX, 7-7

- network communications device, 7-8

- the network communications device

- on VAX, 7-7

- virtual terminals, 7-11

CONSCOPY.COM command procedure, 5-49

Console report during UETP, 17-15

Console storage device

- connecting (VAX), 7-7

- copying, 5-49

- use in booting, 4-2

Console terminals, 2-21, 2-23

message

- indicating lack of installed page file, 5-6
- indicating site-independent startup, 4-5
- indicating site-specific startup, 4-5
- login welcome, 2-10

Container files, 3-20

CONTIN.SYS file

- See Continuation file

- Continuation file
 - CONTIN.SYS, A-9
 - reserved file, A-9
 - used as extension file identifier, A-9
- Continuation volumes
 - in volume set, 8-40
 - mounting in tape volume sets, 8-38
- CONTINUE command
 - in conversational boot, 4-6
- Control access
 - for disk directory files, 9-11
 - for disk files, 9-6
- Conversational boot
 - changing system parameters, 4-6, 14-36
 - CONTINUE command, 4-6
 - location of computer-specific instructions, 4-4
 - performing, 4-3
 - SET command, 4-6
 - SET/STARTUP command, 4-12
 - SHOW command, 4-6
 - showing system parameters, 4-6, 14-36
 - SHOW/STARTUP command, 4-12
 - specifying an alternate startup command
 - procedure, 4-12
 - SYSBOOT prompt, 4-3
 - tasks allowed in, 4-3
 - USE command, 4-7
 - uses of, 4-3
 - with alternate system parameter file, 4-7
 - with minimum startup, 4-14
- Convert utility (CONVERT)
 - saving the queue database, 12-11
 - using to change organization of file, 9-19
- Coordinated Universal Time
 - See UTC service
- COPY command, 10-23, 21-14
 - comparison with COPY command in System
 - Dump Analyzer utility, 15-15
 - disk volumes, 9-19
 - in System Dump Analyzer utility, 15-4, 15-15
 - restriction for copying dump files, 15-15
 - sending message after file is copied, 9-22
 - standard-labeled volumes
 - copying from, 9-20
 - tape volumes
 - copying files from, 9-21
 - copying files to, 9-20, 9-21
 - transferring information, 9-19
- Copying directories
 - with BACKUP, 10-23
- Copying files
 - dump files, 15-14
 - from disk volumes, 9-20
 - from tape volumes, 9-20
 - methods for, 9-19
 - to and from a remote host, 21-14
 - to disk volumes, 9-19
 - to tape volumes, 9-21
- Copying files (cont'd)
 - using COPY command, 9-19
 - using Exchange utility, 9-24
 - with BACKUP, 10-23
- Copying software product release notes, 3-29
- Core image file
 - CORIMG.SYS, A-9
 - not supported by OpenVMS, A-9
- CORIMG.SYS file
 - See Core image file
- Counters
 - status of LAT node, 24-9
- CPUDEFAULT process limit
 - choosing a value for batch queues, 13-31
 - specifying a value for batch queues, 13-19, 13-29
- CPU identification number, 26-2
- CPUMAXIMUM process limit
 - choosing a value for batch queues, 13-31
 - specifying a value for batch queues, 13-19, 13-29
- CPU process limit, 6-42
- CRASH console command, 4-27
- Crash dumps, 15-2
 - See also System dump files
 - See also System failures
 - comparison of physical and selective, 15-3
 - freeing page file of, 15-16
 - physical, 15-3
 - releasing, 15-16
 - requirements for saving, 15-3
 - saving contents of page file on reboot, 15-2
 - saving contents of system dump file on reboot, 5-13
 - selective, 15-3
- Crash Log Utility Extractor (CLUE)
 - description, 15-12
- CREATE command
 - creating directories, 9-20
 - limiting number of file versions, 8-53
 - in SYSGEN
 - changing page, swap, and dump file sizes, 15-26
 - creating page, swap, and dump files, 15-18
 - writing new file to tape volume, 9-18
- CREATE/DIRECTORY command
 - to specify UIC-based directory protection, 9-12
- Creating an additional queue manager, 12-10
- Creating a new system parameter file
 - with SYSGEN, 14-34
- Creating a PCF
 - before installation, 3-22
 - during installation, 3-30
- Creating a queue database, 12-7
- Creating execution queues, 13-15
 - autostart, 13-15, 13-16
 - nonautostart, 13-16

- Creating generic queues, 13-17
- Creating log files
 - operator log file, 18-22
- Creating page, swap, and dump files
 - with AUTOGEN, 15-17, 15-24
 - with SYSGEN, 15-18
- Creating search path of Help Message database files, 5-28
- Cross-architecture booting, 22-23
- CTRLNAME logical name, 17-32
- Current accounting file
 - controlling which resources are tracked in, 19-3
 - default file name, 19-3
 - definition, 19-1
 - finding out what resources are tracked in, 19-2
 - moving, 19-3
 - starting up a new version of, 19-3
- Current system parameters, 14-3, 14-26
- Customizing DECwindows Motif interface, 3-39
- Customizing the system
 - adding optional system files that have been removed from the system disk, 5-1
 - backing up
 - console storage device, 5-49
 - system disk, 5-49
 - building standalone BACKUP, 5-49
 - creating site-specific startup command procedures, 5-2
 - creating systemwide announcements, 5-14
 - duplicating the system disk, 2-26
 - enabling autostart, 5-12
 - installing known images, 5-12
 - installing resident images (Alpha), 5-13
 - limiting the number of interactive users, 5-16
 - making remote InfoServer devices available, 23-14
 - making remote InfoServer disks available, 5-13
 - modifying login procedures, 5-16
 - modifying site-specific startup command procedures, 5-2
 - rules, 5-4
 - SYCONFIG.COM, 5-7
 - SYLOGICALS.COM, 5-8
 - SYPAAGSWPFILES.COM, 5-6
 - SYSECURITY.COM, 5-9
 - SYSTARTUP_VMS.COM, 5-10
 - removing optional system files from the system disk, 5-1
 - setting up a LAT network, 5-15
 - starting InfoServer Client for OpenVMS, 5-13, 23-10
 - starting queues, 5-12
 - starting the DECnet network, 5-15
 - submitting batch jobs at system startup, 5-14

- Cylinders
 - definition, A-2

D

- DAD virtual disk unit, 23-12
- Databases
 - See also Product database configuration
 - See Configuration databases
 - LAT database, 24-8
 - LMF
 - use in system startup, 5-5
 - network
 - permanent, 21-11
 - volatile, 21-11
 - of software products, 3-20
 - queue
 - See Queue database
 - startup
 - definition, 5-18
 - layered product, 5-18, 5-19
 - OpenVMS, 5-18
 - order of startup events, 5-4
- Data blocks
 - partially recorded
 - ISO 9660 standard, 8-5
- Data card deck, 7-18
- Data interleaving
 - ISO 9660, 8-6
- Data loss
 - avoiding by dismounting volume, 8-44
- Date formats, 5-42
 - predefined, 5-47
 - specifying, 5-45
- Daylight saving time
 - changing your time differential factor (TDF) for, 5-35
- DAYLIGHT_SAVINGS.COM command procedure, 5-35
- DBBFs (detected bad block files), 8-65
- DCL commands
 - accessing disk and tape files, 9-13
 - executing with SYSMAN DO command, 2-18
 - file manipulation, 9-1
 - for remote directory listings, 21-14
 - for remote file access, 21-14
 - for remote terminal service, 21-13
 - for system management, 2-5
 - retrieving file information, 9-2
 - support for TCP/IP, 21-13
 - with DO command in SYSMAN, 20-20
- DCL interface, 3-18
- DDCMP (Digital Data Communications Message Protocol)
 - for multipoint connections, 21-9
 - for point-to-point connections, 21-8

ddcu format

- for device names, 7-1

DEALLOCATE command, 8-10

- tape, 9-23

Deallocating devices, 8-10

Deallocating drives, 8-10

DECdns naming service

- storing full names in, 21-2

DECdtm services

- and managing transaction logs, 25-1

- disabling, 25-20

- enabling, 25-21

DECevent utility

- additional commands, 18-11

- description, 18-2

- DIAGNOSE command, 18-2

- exiting, 18-9

- introduction, 18-8

- invoking, 18-9

- privileges required, 18-11

- producing reports, 18-11

- qualifiers, 18-10

- report formats, 18-9

- brief, 18-12

- fast error (FSTERR), 18-15

- full, 18-11

- summary, 18-14

- terse, 18-13

- using to report on error log file, 18-8

DECnet

- See also Networks

- adaptive routing, 21-5

- advantages, 21-1

- asynchronous

- using virtual terminals, 7-11, 7-12

- circuit

- definition, 21-4

- communications line

- definition, 21-4

- configuration

- automatic, 21-12

- manual, 21-12

- on an OpenVMS system, 21-7

- with bridge, 21-7

- configuration database, 21-10

- connecting with communications line, 21-12

- definition, 21-4

- end node, 21-6

- error message during UETP, 17-28

- establishing node in network, 21-11

- Ethernet, 21-5

- Event Logging facility

- to monitor network events, 21-17

- license, 21-12

- local area network (LAN) connections, 21-8

- logical link

- definition, 21-4

- managing a network node, 21-14

DECnet

- managing a network node (cont'd)

- providing host services, 21-15

- managing remote nodes with, 2-12

- multiple-area network, 21-6

- network configurations, 21-6

- network connections, 21-8

- network interface for OpenVMS, 21-7

- network monitoring, 21-14

- See also Networks, monitoring

- tools, 21-15

- nodes

- definition, 21-4

- object definition, 21-4

- overriding AUTOGEN observations, 14-21

- PAK, 21-12

- Phase IV and Phase V compared, 21-1

- planning configuration, 21-12

- preparing for UETP, 17-12

- preparing system, 21-12

- router node, 21-5

- routing

- definition, 21-5

- levels of, 21-5

- security, 21-12

- controlling access to node, 21-12

- protecting files, 21-12

- using proxy accounts, 21-12

- shutting down for software installation, 3-7

- specifying MAIL addresses, 5-33

- terminology, 21-4

- testing network, 21-17

- tools

- DTS/DTR, 21-17

- Monitor utility, 21-17

- NCP Ethernet configurator, 21-17

- UETP defaults for installation, 17-30

- UETP test phase, 17-35, 17-36

- using with EXCHANGE/NETWORK, 9-24

- verifying connection to the network, 21-12

- WAN (wide area network) connections, 21-8

- with VMScluster systems, 21-6

- worldwide connections, 21-9

- DECnet/OSI for OpenVMS, 21-1

- definition, 21-1

- distinguished from DECnet for OpenVMS, 21-1

- full names, 21-2

- starting, 5-16

- DECnet for OpenVMS, 21-1

- node names, 21-2

- starting, 5-15

- DECnet Test Sender/DECnet Test Receiver utility (DTS/DTR)

- network monitoring tool, 21-17

- DECW\$TAILOR

- See Tailoring utilities

- DECwindows
 - use of Snapshot facility with, 4-21
- DECwindows Motif interface, 3-18
 - customizing your environment, 3-39
 - POLYCENTER Software Installation utility, 3-34
 - buttons, 3-37, 3-39, 3-40
- Deductible resource, 6-3
- DEFAULT account
 - in UAF, 6-6
- Default boot procedure, 4-2
- Default directories, 6-15
- Default form, 13-63
- DEFAULT form, 13-63
- Default protection
 - UIC, 9-7
- Default system parameters
 - booting with, 4-8
- DEFINE/CHARACTERISTIC command, 13-58
- DEFINE command
 - NCP, 21-11
- DEFINE/FORM command, 13-61
 - for controlling line overflow, 13-44
- Defining a form, 13-61
- Defragmenting disks, 10-48
- Deinstalling layered software products, 3-34
- Delete access
 - explicitly assigning, 9-11
 - for disk directory files, 9-11
 - for disk files, 9-6
 - granting through protection codes, 11-8
- DELETE/CHARACTERISTIC command, 13-59
- DELETE/ENTRY command, 13-81
- DELETE/FORM command, 13-64
- DELETE/QUEUE command, 13-57
- Deleting
 - Digital messages from the Help Message database, 5-29
 - files from the system disk, 5-1
 - forms, 13-64
 - problems with, 13-82
 - jobs, 13-74
 - page, swap, and dump files
 - after creating new version, 15-28
 - caution, 15-21
 - queue characteristics, 13-59
 - problems with, 13-82
 - queues, 13-57
 - problems with, 13-82
- Dependent software products, 3-27
- DESELECT command
 - in SHOW CLUSTER, 20-13
- Despooling a spooled printer, 7-15
- Destination for installing software
 - copying files, 3-33
 - PCSI\$DESTINATION location, 3-28
 - specifying location, 3-28
- Destination parameter
 - in VMSINSTAL.COM, 3-12
- Detected bad block files (DBBFs)
 - See DBBFs
- Device control libraries, 13-44 to 13-45
 - See also Device control modules
 - assigning to a queue, 13-66
 - order of module output, 13-45
 - procedure for using, 13-65
 - sample commands, 13-68
 - setting up, 13-44
- Device control modules, 13-44, 13-83
 - See also Device control libraries
 - adding, 13-66, 13-83
 - creating, 13-66
 - deleting, 13-66, 13-83
 - forms, 13-65, 13-67
 - inserting into a library, 13-66
 - listing, 13-67
 - naming, 13-67
 - order of output, 13-45
 - page setup, 13-44, 13-67
 - requesting with PRINT command, 13-67
 - reset, 13-67
 - when queue is started, 13-67
 - sample commands, 13-68
 - setting up, 13-65
 - setup, 13-44, 13-67
 - specifying, 13-45, 13-65
 - storing, 13-66
 - troubleshooting, 13-83
 - types, 13-44
 - with forms, 13-45
- Device drivers
 - CONINTERR.EXE file, 7-7
 - for event handling, 7-7
 - loading
 - automatically, 7-6
 - in system startup, 5-4, 5-7
 - manually (Alpha), 7-8
 - manually (VAX), 7-7
 - not associated with a specific device, 7-7
 - TTDRIVER, 7-11
- Device names, 7-1
 - for virtual terminals, 7-11
 - in a VMScluster environment, 7-1
- DEVICE phase of system startup, 5-4, 5-18
- Devices
 - accessing in batch job, 8-46
 - allocating, 8-9, 9-22
 - availability
 - OPCOM message, 8-59
 - configuring
 - in system startup, 5-7, 7-6
 - manually, 5-7, 7-7
 - special devices, 7-7, 7-8
 - connecting, 7-6

- Devices (cont'd)
 - determining available, 7-2
 - Ethernet adapter
 - specifying number for AUTOGEN, 14-21
 - getting information about, 7-2, 7-16
 - ISO-9660
 - getting information about, 7-5
 - LTA_n, 7-11
 - magnetic tape
 - See Tapes
 - managing, 7-1
 - managing with LAN Control Program (LANCP)
 - utility, 22-9
 - manually configuring non-standard devices, 5-7, 7-7
 - mounting volumes, 8-18
 - network communications
 - connecting, 7-8
 - connecting (VAX), 7-7
 - not recognized by the system, 7-7
 - OPA0:, 2-23
 - printers
 - See Printers
 - protecting, 7-5
 - requiring manual connecting, 7-8
 - resetting default in SYSMAN, 2-17
 - RTA_n, 7-11
 - setting characteristics, 13-14
 - in system startup, 13-14
 - special
 - connecting, 7-8
 - spooled, 13-14
 - status report on, 18-17
 - suppressing autoconfiguration during system startup, 5-8, 7-8
 - terminals
 - See Terminals
- Device tests
 - running individually with UETP, 17-32
- Device unavailable queue status, 13-51
- DIAGNOSE command, 18-2
- Diagnostics
 - relationship to UETP, 17-16
- Dialup identifiers, 11-11
- DIBOL
 - starting the message manager, 5-16
- Digital Data Communications Message Protocol
 - See DDCMP
- Digital System Identifier
 - See DSI
- Digital writers
 - sending comments to, iii
- DIOLM process limit, 6-42
 - value for efficient backups, 10-10
- Direct I/O count process limit, 6-42
- Direct mode
 - as execution mode for a startup procedure, 5-19
- Directories
 - backing up, 10-24
 - backlink, 8-56
 - copying with BACKUP, 10-23
 - creating, 9-20
 - destination
 - specifying in VMSINSTAL.COM command procedure, 3-13
 - for an interactive account, 6-15
 - levels of access in restore operations, 10-30
 - protecting, 9-11, 9-12
 - resetting default in SYSMAN, 2-17
 - temporary working
 - in VMSINSTAL.COM command procedure, 3-14
- DIRECTORY command, 9-5
 - checking number of user's disk blocks, 8-49
 - retrieving file information, 9-2
 - showing full information, 9-17
 - to obtain product list, 3-10
 - with tapes, 9-16, 9-20
- Directory trees
 - copying, 10-23
- DIR/FTP command, 21-14
- DISABLE AUTOSTART/QUEUES command, 13-54, 13-55
 - entering before shutting down a system, 13-55
 - relationship to STOP/QUEUES/ON_NODE command, 13-55
- Disabled queues, 13-84
- Disabling autostart, 12-9, 13-54
 - before shutting down a node, 13-55
- Disabling operator terminals, 2-24
- Disabling spooling, 7-15
- Disabling user accounts, 6-27
- Disk
 - See System disks
 - See User disks
- Disk files
 - accessing at file level, 9-13
 - assigning an alias, 9-10
 - copying
 - from disk volumes, 9-19
 - to tapes, 9-21
 - using COPY command, 9-18
 - modifying characteristics, 9-8, 9-10
- DISKQUOTA commands
 - See also Disk Quota utility
- DISKQUOTA commands in SYSMAN, 2-11
- DISKQUOTA/DISABLE command, 8-52
- DISKQUOTA/ENABLE command, 8-52

Disk quotas, 6-15

See also Disk Quota utility

creating, 8-51

definition, 8-48

disabling, 8-53

displaying, 8-49

ensuring accuracy with rebuild operation, 8-50

establishing, 8-50, 8-51

example, 6-33

exceeding, 8-50

file, 8-49

maintaining, 8-50

retrieving information, 8-52

suspending operations, 8-52

Disk Quota utility (DISKQUOTA), 8-50

Disks

See also CD-ROMs

See also Disk commands

See also Disk files

See also Disk quotas

See also Disk volumes

See also System disks

allocating drives, 8-9

allocating space on, 8-52

backing up active, 10-63

block

definition, 8-2

grouped into cluster, A-1

changing volume label, 3-33

cluster, A-1

definition, 8-2

concepts, A-1

cylinder

definition, A-2

deallocating drives, 8-10

default format, 9-19

definition, 8-2

dismounting, 10-16

extents, A-1

definition, 8-2

file identification, A-4

files

See Disk files

Files-11

directory hierarchy, A-4

fragmentation of, 10-48, 15-27

I/O performance, 16-8

initializing, 10-14

mounting, 8-28, 10-15

mounting in host-based shadow set, 10-42

organization

logical, A-1

physical, A-2

protecting, 8-14, 8-16

space

See Disk space

structure

Disks

structure (cont'd)

See Disk structure

system

See System disks

terminology, 8-2

track

definition, A-2

usage, 8-49

creating file, 8-58

UICs kept in quota file, A-9

volumes

definitions, 8-2

Disk space

See also Disk quotas

allocation by cluster, A-1

managing, 8-48 to 8-55

purging files, 8-53

saving, 8-53

by moving page and swap files off the
system disk, 15-5

by purging the operator log file, 5-14

by removing optional system files, 5-1

by storing minimal dump information,
15-10

by using a selective dump, 15-3, 15-4

tracking use of, 19-6

Disk storage server, 23-10

Disk structure

analyzing errors, 8-55

Files-11, A-3

reporting and repairing errors, 8-56

Disk volumes

accessing files on, 9-13

access to public, 8-8

adding to an existing set, 8-34

adding volumes to volume sets, 8-34

analyzing disk structure errors, 8-55

analyzing media errors, 8-64

assigning logical name, 5-11

assigning volume label, 5-11

binding into volume sets, 8-30

characteristics

modifying, 8-29

console, 9-24

copying files from, 9-20

copying files to and from foreign volumes, 9-24

copying files to tape volumes, 9-22

creating Files-11 structure, 8-10

creating volume sets from, 8-32

definition, 8-2, A-1

disk quota operations, 8-50

dismounting, 8-43

file expiration dates

setting, 8-54

file-structured, 8-20

foreign, 8-20

handling error conditions, 8-55

Disk volumes (cont'd)

- initializing, 8-10, A-1
 - guidelines, 8-13
- load balancing, 8-9
- modifying characteristics, 8-29
- mounting, 8-21, 8-27
 - for page and swap files, 15-20
 - for queue database files, 12-6
 - in system startup, 5-11
 - early, 5-12
 - for page and swap files, 5-6
 - InfoServer, 23-14
 - MOUNT/ASSIST command, 5-11
 - special consideration about operator assistance, 5-11
- mount verification, 8-58
- performance, 8-28
- physical loading, 8-21
- private, 8-9
- protecting, 8-14
- public
 - See Public volumes
- reading files from, 9-13
- rebuilding, 8-50
- removing before dismounting, 8-43
- space

- conserving, 8-48

- verification, 8-58

- write-locked

- dismounting, 8-43

- write-locking, 8-60

- writing files to, 9-20

DISMOUNT command

- See also Dismounting

- canceling mount verification, 8-62

- dismounting single volume in volume set, 8-44

- for a single tape volume, 8-44

- for foreign volumes, 8-45

- for single volume or volume set, 8-42

- overriding automatic unloading of volume, 8-45

- tape, 9-23

Dismounting

- See also DISMOUNT command

- a backup volume, 10-16

- a disk with DECdtm transaction logs, 25-14

- system disk, 8-43

- volumes

- conditions preventing, 8-42

- foreign, 8-45

- in a VMScluster system, 8-45

- with cached information, 8-43

- with open files, 8-43

- volume sets, 8-45

Displaying

- characteristics assigned to a queue, 13-59

- defined characteristics, 13-59

Displaying (cont'd)

- defined forms, 13-63

- forms assigned to queues, 13-64

- information about queues, 13-49

- information about the queue manager, 12-11

- system parameters

- See Showing system parameters

Distributed Queuing System

- See DQS

- Distributing system work load, 16-4

Distribution kit

- startup files included on, 5-2

DO command

- for managing a VMScluster system, 20-20

- in SYSMAN, 2-18

Documentation

- sending comments to Digital writers, iii

DOS-11 tape volumes

- file transfers with, 9-24

- format conversions for, 9-24

Downline loading, 23-2

DQS (Distributed Queuing System)

- distributed printing, 13-13

Drivers

- See Device drivers

- DS3 connection, 21-9

- DSA device naming, 7-1

DSI (Digital System Identifier)

- ISO 9660 media protection, 8-18

- mount option, 8-17

DSI keyword

- with MOUNT/PROTECTION command, 8-23

DTS/DTR

- See DECnet Test Sender/DECnet Test Receiver utility

Dual-architecture VMScluster systems

- installing images, 20-20

- example, 20-21

- DUMPBUG system parameter, 15-3

Dump file information

- saving automatically, 15-11

- Dump file off system disk, 15-14

- requirements, 15-14

Dump files

- changing sizes

- with SWAPFILES.COM, 15-25

- system

- See System dump files

- DUMPFIL symbol, 15-24

- DUMPSTYLE system parameter, 15-3, 15-7, 15-10

- Dynamic load balancing, 26-2

- Dynamic system parameters, 14-2, 14-3

- See also System parameters

E

EDIT keypad function, 20-7

Emergency system shutdown

with console commands, 4-27

with OPCCRASH, 4-27, 4-36

Emergency system startup

with default system parameters, 4-8

without startup and login procedures, 4-8

without the UAF, 4-10

ENABLE AUTOSTART/QUEUES command,
13-18, 13-48

in startup command procedure, 13-16

recommended use, 13-18

Enabling autostart, 13-4, 13-18, 13-48

in startup command procedure, 13-16

End nodes

network, 21-6

END phase

of system startup, 5-19

ENQLM process limit, 6-42

EOT markers, 8-6

continuing to copy after, 9-24

Erasing blocks, 8-58

ERLBUFFERPAGES system parameter, 15-6

ERRFMT

See Error Formatter

ERRFMT process, 18-2

See also Error log files

See also Error logging

See also Error Log utility

creation during system startup, 5-5

restarting, 18-3

sending mail when it is deleted, 18-4

writes to ERRLOG.SYS file, 18-2

Error checking

in SYSTARTUP_VMS.COM, 5-10

in system parameter files, 14-19

Error Formatter

changing mail username, 18-5

description, 18-2

disabling mail, 18-4, 18-5

enabling mail, 18-5

notifying user with mail message, 18-4

using to send mail, 18-4

ERRORLOGBUFFERS system parameter, 15-6

Error log files

created by ERRFMT process, 18-2

events reported in, 18-5

logical name defining location, 5-8

maintaining, 18-3

moving to reduce system disk I/O, 16-8

producing full report, 18-6

reporting on using DECEvent, 18-8

SYSPRV privilege to access, 18-6

Error logging

See also ERRFMT process

See also Error log files

See also Error Log utility

description, 18-2

reports produced, 18-3

using the Error Log utility, 18-2

Error logging facility, 18-2

Error Log Report Formatter (ERF), 18-2

Error log reports

See Error Log utility, reports

Error Log utility (ERROR LOG)

See also ERRFMT process

See also Error log files

See also Error logging

ANALYZE/ERROR_LOG command, 18-6

definition, 18-5

relationship to UETP, 17-3, 17-17, 17-29

reporting on error log file, 18-5

reports

excluding unknown entries, 18-8

formats, 18-7

printing, 18-7

privileges required, 18-7

specifying display device, 18-8

specifying events and times, 18-8

uses, 18-5

report types, 18-5

Error messages, 2-5

See also Help Message utility

See also Messages

Error options

for fatal BACKUP errors, 10-64

Errors

analyzing disk structure, 8-55

analyzing error reports, 18-2

analyzing media, 8-64

disk read

if returned when booting, 4-16

disk structure

repairing, 8-57

reporting, 8-57

error log file, 18-2

error logging facility, 18-2

handling on disk volumes, 8-55

machine check

if returned when booting, 4-16

mounting disk, 8-22

Errors during UETP, 17-21

diagnosing, 17-16

sources of, 17-17

ESS\$LASTDRIVER device driver, 23-9, 23-13

controlling and diagnosing, 23-9, 23-10

- ESS\$STARTUP.COM command procedure, 23-10, 23-13
 - invoking in system startup, 23-10
- Ethernet
 - adapters
 - specifying number of in MODPARAMS.DAT file, 14-21
 - configurator
 - network monitoring tool, 21-17
 - connecting, 21-7
 - defining a remote node for UETP, 17-19
 - definition, 21-5
 - device drivers, 22-2
 - in local area network, 21-8
 - linking with bridge, 21-7
 - preparing for UETP, 17-8
- Event classes
 - displaying those being audited, 18-27
- Event handling
 - device driver used in, 7-7
- EXCHANGE/NETWORK command
 - using to transfer files, 9-19, 9-24
- Exchange utility (EXCHANGE)
 - using to copy files, 9-19
 - using to transfer information, 9-24
- Executable images, 16-9, 16-14, 16-15
- Execute access
 - for disk directory files, 9-11
 - for disk files, 9-6
 - granting through protection codes, 11-8
- Execute procedure (@) command, 20-16
- Executing job status, 13-69, 13-74
- Execution modes
 - startup procedures, 5-19
 - BATCH, 5-19
 - changing, 5-21
 - DIRECT, 5-19
 - SPAWN, 5-19
 - specifying, 5-20
- Execution queues
 - activating autostart, 13-4, 13-15, 13-49
 - creating, 13-15
 - relationship to generic queues, 13-2
 - starting
 - autostart, 13-4, 13-18, 13-48
 - in system startup, 5-12
 - nonautostart, 13-16, 13-18, 13-48
- Executive mode
 - calling images running in, 16-11, 16-14
 - logical names, 16-15
 - recommended use for logical names, 5-9
- Expiration date, 6-42
 - field, 9-14
 - checking, 9-18
 - file, 8-54
 - tape file system checks, 9-17

- Expiration time, 6-42
- Extended attribute records
 - See XARs

- Extensions
 - file
 - See File extensions

- Extents
 - definition, A-1
 - disk
 - definition, 8-2
 - index file
 - definition, A-7

- Extracting software product release notes, 3-29

F

- F\$GETJPI lexical function
 - getting information about vector processing, 26-9
- F\$GETQUI lexical function, 13-51
- F\$GETSYI lexical function
 - getting information about vector processing, 26-9
- Failover list
 - for an autostart queue
 - specifying, 13-15
 - for queue manager, 12-3, 12-9
 - insufficient, 12-17
 - specifying, 12-9
- Failovers
 - See also Failover list
 - of queue manager, 12-3, 12-17
 - forcing, 12-9
 - of queues, 13-4, 13-15
- FDDI (Fiber Distributed Data Interface)
 - device drivers, 22-2
 - multiaccess communications channel, 21-8
 - supports VAXcluster technology, 21-8
- Feedback
 - See AUTOGEN feedback
- Feedback on documentation
 - sending comments to Digital writers, iii
- Fiber Distributed Data Interface
 - See FDDI
- FID (file identification)
 - See File identification
- FIELD account
 - initial modification, 6-9
 - in UAF, 6-7
- Field of data
 - in SHOW CLUSTER, 20-4
 - removing, 20-11
- Field Service accounts
 - in UAF, 6-7
- File access
 - disk, 9-13
 - levels allowed in restore operations, 10-30

File access (cont'd)

- listing number of concurrent, 16-8
- of a remote host, 21-14
- tape, 9-13, 9-14

File banner pages, 13-42, 13-60

- See also Job banner pages

File extensions

- effect on system performance, 16-7
- specifying size, 8-22, 16-7
- system parameter controlling, 16-7

File formats

- use with BACKUP, 10-9

File fragmentation

- of page and swap files, 15-27

File headers

- index file, A-7
 - contents, A-7
 - extension, A-8
 - primary, A-8

File identification

- file number, A-4
- Files-11, A-4
- file sequence number (SEQ), A-4
- relative volume number (RVN), A-4

File Log

- VMSINSTAL.COM option, 3-17

File names

- OpenVMS extended, 9-16
- standard, 9-16

File protection, 9-4

- SYSDUMP.DMP file, 15-4

Files

- See also Files-11 On-Disk Structure
- See also Parameter files
- accessing
 - See File access
- attributes
 - accessing, 9-17
- backing up, 10-23
- comparing with BACKUP, 10-25
- copying
 - from disk to standard-labeled volumes, 9-20
 - from disk volumes, 9-19
 - remotely, 21-14
 - to tape, 9-21
 - to tape volumes, 9-21
- copying with BACKUP, 10-23
- creating, 8-3
- detected bad block (DBBF), 8-65
- expiration dates on, 8-54
- for AUTOGEN feedback, 14-11
- limiting number of versions, 8-53
- logging activity during installation, 3-17
- lost
 - recovering, 8-57
- modifying characteristics, 9-8
- naming

Files

naming (cont'd)

- on Files-11 volume, A-7

nonstandard format

- DCL commands with, 9-2, 9-13

on public volumes, 8-8

open during backup, 10-30, 10-63

overwriting, 9-14

PCF, 3-22

private volumes, 8-9

privileges, 9-4

public, 8-8

purging to save disk space, 8-53

recovering lost, 8-57

remote access, 21-14

reserved, A-4

- list of, 8-4

restoring from image backup, 10-42

restoring from incremental backup, 10-44

restoring with BACKUP, 10-28

retrieving information from, 9-2

security

- using protection codes, 9-6

system

- moving to reduce system disk I/O, 16-8

tape

- See Tape files

tape volumes, 9-5

- writing to files on, 9-18

transferring across network, 9-24

- using COPY, 21-14

versions

- limiting number of, 8-53, 9-12

VMSMAIL_PROFILE.DATA, 6-38

Files-11 On-Disk Structure

block

- definition, A-1

CD on OpenVMS, 8-5

comparison of ODS Level 1 and Level 2, A-10

creating structure, 8-10

disk save set, 10-7

file identification, A-4

master file directory (MFD), A-4

ODS Level 1, 8-3

- directory hierarchy, A-4

ODS Level 2, 8-3

- assigning disk quotas, 8-50

- definition, 8-4

- directory hierarchy, A-4

sector, A-2

structure, 8-3, A-3

- Level 1, 9-19, A-4

- Level 2, 9-19, A-4

- reserved files, 8-4

terminology, A-2

UIC, A-4

using with Exchange utility (EXCHANGE)

- to transfer data, 9-24

- File specifications
 - ANSI, 9-16
 - for installing images, 16-16
- File structures, 8-3
- File Transfer Protocol
 - See FTP
- File windows
 - mapping pointers for, 8-24
- FILLM process limit, 6-43
 - value for efficient backups, 10-10
- Firmware
 - LAN updates, 22-14
- Flag pages, 13-33
 - file, 13-35
 - job, 13-33
- Foreign volumes
 - See also Volumes, foreign
- Formats
 - CD-ROM on disc, 8-4
 - converting software kits, 3-33
 - High Sierra, 8-4
- Formatting volumes, 8-11
- Form feeds
 - controlling page overflow with, 13-44
 - inserting automatically in output jobs, 13-19
- Forms
 - assigning default for a queue, 13-63
 - associating with jobs and queues, 13-42
 - changing DEFAULT, 13-63
 - commands used with, 13-61
 - controlling line overflow with, 13-42, 13-43
 - controlling page width, length and margin size with, 13-42
 - controlling paper stock with, 13-42
 - default, 13-63
 - DEFAULT, 13-63
 - defining, 13-61
 - deleting, 13-64
 - problems, 13-82
 - description, 13-42
 - displaying
 - defined forms, 13-63
 - forms assigned to queues, 13-64
 - formatting jobs with, 13-42
 - mounting, 13-64
 - preprinted
 - aligning, 13-75, 13-77
 - procedure for using, 13-43, 13-60
 - specifying setup modules with, 13-42
 - specifying sheet-feed paper with, 13-42
- Fragmentation of disks, 10-48
- FTP (File Transfer Protocol), 21-14
- Full backup
 - See Image backup
- Full names
 - DECnet/OSI
 - assigning, 21-3

- Full names
 - DECnet/OSI (cont'd)
 - syntax, 21-2
 - definition, 21-2

G

- GBLPAGES system parameter, 16-12
- GBLSECTIONS system parameter, 16-12
- Generic queues
 - batch, 13-3
 - recommended use, 13-7
 - creating, 13-17
 - description, 13-2, 13-3
 - in a VMScluster environment, 13-2
 - output, 13-3, 13-11, 13-12
 - recommended use, 13-11
 - relationship to execution queues, 13-2
 - specifying target execution queues, 13-17
- Get Save Set
 - VMSINSTAL.COM option, 3-15, 3-16
- Getting help, 3-30
- GHRs (granularity hint regions)
 - slicing shareable images, 16-12
- Global pages, 16-12
- Global sections, 16-12
- Granularity hint regions
 - See GHRs
- Group numbers
 - modifying, 20-16
- Groups
 - accounting, 19-5
- Group users (security category), 11-8
- Group volumes
 - definition, 8-8
- GRPPRV privilege
 - giving rights of system user, 11-8

H

- Halting system
 - waiting until after system dump file is written, 15-2
- Hardware
 - booting problem, 4-16
 - importance of sufficient capacity for system performance, 16-6
- Header labels
 - on tape files, 8-7, 9-15
 - reading attributes of, 9-17
- Header resident images, 16-10, 16-12
- Help
 - POLYCENTER Software Installation utility, 3-30, 3-35
- Help Message utility (MSGHLP), 5-25
 - accessing \$STATUS values for uninstalled messages, 5-26

Help Message utility (MSGHLP) (cont'd)

- adding .MSGHLP\$DATA files to the database, 5-28
- adding comments to the database, 5-30
- adding messages to the database, 5-32
- changing Digital-supplied data, 5-31
- compressing the database after deletions, 5-30
- creating databases for different user groups, 5-28
- customizing the database, 5-25
- default database, 5-28
- deleting Digital-supplied messages from the database, 5-29
- message section files, 5-26
 - order searched, 5-26
- search path of database files, 5-28
- /SECTION_FILE qualifier usage, 5-26

Hierarchical storage controller devices

See HSC devices

High Sierra format

- definition, 8-4
- on CD-ROM, 8-4

Highwater marking

- disabling for system performance, 16-7

Holding a job, 13-71

Holding job status, 13-69, 13-79

Home blocks, 8-58

- in index file, A-6

Host-based shadow set

- mounting disks, 10-42

HSC devices

- configuring during system startup, 7-6
- disabling configuration during system startup, 7-8

Hyphen (-), restriction on use of, 3-18

I

I/O (input/output)

- reducing on system disk, 16-8

Identification record

- ANALYZE/DISK_STRUCTURE, 8-56

Identifier field

- file, 9-15
- volume, 8-39

Identifiers

- Environmental, 11-10
- General, 11-10
- system-defined, 11-10
- types, 11-10
- UIC, 11-10

Idle queue status, 13-51

Image backup

- command format for disks, 10-32
- command format for tapes, 10-31
- definition, 10-2
- restoring files from, 10-42
- to disk, 10-32

Image backup (cont'd)

- to tape, 10-31, 10-32

Image Registry facility, 5-22

Images

See also Known images

- concurrent access by multiple users, 16-11
- definition, 16-9
- determining frequency of use, 16-15, 16-16
- executable, 16-9, 16-14, 16-15
- execute-only, 16-15
- header resident, 16-10, 16-12
- installing
 - See also Known images
 - effect on RUN command, 16-10
 - in system startup, 5-4, 5-12, 5-13, 16-10
 - reasons for, 16-9
 - to improve system performance, 16-7
- known, 16-10

See also Known images

linkable, 16-9

permanently open, 16-10, 16-12

privileged, 16-10, 16-13, 16-14

security caution, 16-13

privileged shareable, 16-14

protected, 16-11, 16-14

protecting installed, 16-15

relinking to improve system performance, 16-7

resident (Alpha), 16-11

running in protected modes, 16-11, 16-14

shareable, 16-11, 16-14

assigning logical names for, 16-17

slicing on Alpha systems, 16-12

system version dependent

registering, 5-22

user-level

calling of protected code, 16-11, 16-14

version checking, 5-23

writable, 16-11

Incremental backup

- command format for disks, 10-35
- command format for tapes, 10-34
- definition, 10-2
- restoring files from, 10-44
- to disk, 10-34
- to tape, 10-33

INDEXF.SYS file

See Index files

Index files, A-4 to A-8

- alternate file header, A-6
- backup home block, A-5
- backup index file header, A-6
- bitmap, A-6
- boot block, A-5, A-6
- bootstrap image, A-6
- definition, A-5
- file headers, A-6, A-8
- file number, A-4

Index files (cont'd)

- home block, A-5, A-6
- INDEXF.SYS, A-5
- in volume sets, 8-32
- reserved file, A-5

InfoServer

See also InfoServer Client for OpenVMS

- automatic service, 23-4
- availability, 23-4
- backing up system disks, 10-60

Client

- and DECnet, 23-10
- commands, 23-7
- console terminal, 23-6
- determining default server name, 23-6
- downline loading with, 23-2
- fail over, 23-4
- functions, 23-1
- Help facility, 23-8
- load balancing, 23-6
- local connections, 23-6
- management session, ending, 23-7
- mounting devices
 - in system startup, 23-14
- multicast address feature, 23-5
- protocols, 23-5
- quick access to duplicate services for clients, 23-4
- relationship to client systems, 23-2
- remote connections, 23-6
- removing media, 23-4
- server name, 23-6
- service disconnection, 23-4
- setting up in system startup, 23-10
- software
 - starting client for OpenVMS, 23-10
- starting a session, 23-6
- starting Client, 23-10
- support for X terminals, 23-4
- system overview, 23-1
- virtual device server, 23-1
- virtual device units, 23-13
- X terminal client, 23-4

InfoServer Client for OpenVMS

- components, 23-8
- functions, 23-8
- mounting devices
 - in system startup, 5-13
- setting up in system startup, 5-13
- software, 23-9
 - starting automatically, 23-10
 - startup restrictions, 23-12

InfoServer password, 23-6

Initialization files

- creating, 20-15
- establishing SHOW CLUSTER reports, 20-5
- SHOW CLUSTER
 - creating, 20-14

Initialization files (cont'd)

- SHOW_CLUSTER\$INIT, 20-14, 20-15
- use with SYSMAN, 2-19

Initialization of system

- in a multiprocessing system, 26-2

INITIALIZE command

See also Disk commands

See also INITIALIZE/QUEUE command

See also Initializing, volumes

- creating volume identifiers for continuation
 - volumes, 8-39
- disk volumes, 8-11
- for formatting page and swap file disks during
 - system startup, 5-6
- mounting volume sets, 8-39
- qualifiers, 8-12
- setting device protection, 8-16
 - tape volumes, 9-12, 9-22
- tape volumes, 9-21, 9-22
- to format and write label to volume, 8-11

INITIALIZE/QUEUE command, 13-16, 13-48

- activating autostart queues, 13-15, 13-49
- assigning a default form, 13-63
- assigning characteristics, 13-59
- canceling characteristics, 13-59
- controlling page overflow, 13-44
- creating a generic queue, 13-17
- mounting a form, 13-64
- setting block limits, 13-32
- setting scheduling policy, 13-32
- setting UIC-based protection on queues, 13-23
- specifying autostart information, 13-15
- specifying banner pages, 13-60
- specifying job processing options, 13-31
- specifying queue options, 13-19
- specifying reset modules, 13-65
- starting a nonautostart queue, 13-16

Initializing

- queues, 13-15

See also INITIALIZE/QUEUE command

volumes

See also Disk commands

See also INITIALIZE command

- assisting users, 8-14
- disk volumes, 8-10, 8-13
- results of, 10-13
- tape volumes, 9-5

INITIAL phase of system startup, 5-4, 5-18

Initial System Load

See ISL

Input an existing PCF, 3-30

Input specifier

- to the BACKUP command, 10-4

Input symbiont for card reader

- running interactively, 7-18

Installation procedures

See also Installing software

See also POLYCENTER Software Installation utility

See also VMSINSTAL.COM command procedure completing, 3-13

definition, 3-4

on Alpha systems, 3-2

to run VAX system as a C2 system, 3-5

INSTALL command

in SYSGEN, 15-19

in system startup, 5-7

Installed files

See also Known images

See Known images

Installing images, 16-16

See also Install utility

See also Known images

effect on RUN command, 16-10

in SYSTARTUP_VMS.COM, 5-12, 5-13, 16-10

reasons for, 16-9

to improve system performance, 16-7, 16-10

Installing page and swap files

in system startup, 5-6, 15-5, 15-19, 15-21

with AUTOGEN, 15-18, 15-22

with SYPAGSWPFILES.COM command

procedure, 15-19

with SYSGEN, 15-19

Installing software

See also Installation procedure

See also POLYCENTER Software Installation utility

See also Software products

See also VMSINSTAL.COM command procedure as a batch job, 3-32

completing procedure, 3-13

creating a new PCF during, 3-30

deinstalling software, 3-34

identifying kit location, 3-28

logging file activity, 3-17

managing installed software, 3-34

multiple products at once, 3-29

noncompliant products, 3-27

on alternate disk, 3-12

preparing for installation, 3-6

reconfiguring product options, 3-32

removing installed software, 3-34

responding to questions during installation, 3-30

shutting down DECnet, 3-7

specifying source and destination locations, 3-28

supplying answers from a PCF, 3-30

tracking dependencies, 3-27

using a PCF for consistency, 3-22

Install utility (INSTALL)

See also Installing images

See also Known images

and version checking, 5-23

determining number frequency of image use, 16-15, 16-16

improving system performance, 16-7, 16-9, 16-10

listing number of concurrent file accesses, 16-8

making images header resident, 16-10, 16-12

making images privileged, 16-10, 16-14

permanently open images, 16-10, 16-12

reasons for using, 16-9

/RESIDENT qualifier on Alpha, 16-13

slicing feature on Alpha, 16-13

Interactive identifiers, 11-11

Interactive users

limiting for performance management, 16-4

limiting in system startup, 5-16

Interchange environment

protection, 8-20

Interfaces to POLYCENTER Software Installation utility

DCL, 3-18

DECwindows Motif, 3-18

Interfaces to the Backup utility, 10-4

Internet, 21-13

Intrusion databases, 11-7

Intrusions

detection and evasion, 11-6

Intrusion services

managed by security server, 5-5

Invoking AUTOGEN, 14-9

Invoking SDA automatically, 15-11

IO AUTOCONFIGURE command

in SYSMAN, 7-6, 7-8

in system startup, 5-4, 5-7, 7-6

IO CONNECT command

in SYSMAN, 7-8

in system startup, 5-7

IO LOAD command

in SYSMAN, 7-8

IPC (Interrupt Priority C)

using to cancel mount verification, 8-63

IRG (interrecord gap), 8-6

ISL (Initial System Load), 23-2

ISO 9660 standard

data interleaving, 8-6

establishing default file attributes, 8-24

format

definition, 8-4

groups

mounting, 8-34

media

showing device information, 7-5

media protection, 8-17

mounting a volume for, 8-23

ISO 9660 standard (cont'd)

- partially mounted volume sets, 8-35
- partially recorded data blocks, 8-5
- restrictions, 8-36
- standard on OpenVMS, 8-5
- UNDEFINED record format errors, 8-36
- volume label and volume set label duplication, 8-36
- volume labels, 8-36
- volume set labels, 8-36
- volume sets
 - mounting, 8-34
 - partially mounted, 8-35
- volume structure, 8-3

J

Job banner pages, 13-42, 13-60

See also File banner pages

\$JOB card, 7-17

Job controllers

See also JOBCTL process

See also Queue manager

and batch jobs, 12-3, 13-2

communication with queue manager, 12-3

relationship to queue manager, 12-3, 12-7

starting queue manager, 12-6, 12-7

tasks performed by, 12-3

JOBCTL process

See also Job controllers

creation during system startup, 5-5

Job retention

changing for a job, 13-74

specifying for a job, 13-26

specifying for a queue, 13-26

Jobs

See also Batch jobs

See also Output jobs

changing scheduling priority, 13-72

controlling print position and alignment,
13-75, 13-77

deleting, 13-74

holding, 13-71, 13-79

merging, 13-56

modifying, 13-70

moving from one queue to another, 13-57

releasing, 13-71

requeuing

executing, 13-73

pending, 13-73

retaining in a queue, 13-74

suspending, 13-75

Job scheduling, 13-32, 13-72

for output jobs, 13-32

priority

See Priority, job scheduling

Job status

definitions, 13-69

error, 13-26, 13-31, 13-54, 13-74

holding, 13-71, 13-79

pending, 13-73, 13-74, 13-79

retained, 13-71, 13-74

showing, 13-26, 13-68

use in job retention, 13-26

Job table quota, 6-43

Journal file of backup information, 10-26

listing contents of, 10-27

Journal file of queue database, 12-4

See also Queue database

location, 12-6, 12-7

changing, 12-6

JTQUOTA process quota, 6-43

K

Kernel mode

calling images running in, 16-11, 16-14

logical names, 16-15

Keypad definitions, 20-5, 20-7

Kits

See Software products

Known file lists

definition, 16-10

in system startup, 5-12

Known images

definition, 16-10

deleting, 16-18

dismounting volume, 16-18

displaying, 16-16

evaluating merits of installing, 16-15, 16-16

file specification for, 16-16

installing, 16-16

in system startup, 5-12, 16-10

privilege enhancement, 16-13

removing, 16-18

resident, 16-12

L

Labels

backup tape, 10-21

changing volume, 3-33

header, 8-24

initializing volume to write, 8-11

trailer, 8-7

LAD Control Program (LADCP) utility, 23-13 to 23-14

BIND command, 23-14

exiting, 23-13

Help facility, 23-14

invoking, 23-13

making remote InfoServer devices available

locally, 23-14

summary, 23-13

LADCP (LAD Control Program)

See LAD Control Program utility

LANACP utility

See LAN Auxiliary Control Program utility

LAN Auxiliary Control Program (LANACP) utility

OPCOM messages displayed, 22-22

running, 22-6

stopping, 22-6

LAN Control Program (LANACP) utility

servers, 22-5

LAN Control Program (LANCP) utility, 22-6

clearing counters, 22-25

deleting device information, 22-17

deleting node information, 22-20

device database management, 22-14

device management, 22-9

disabling MOP downline load service, 22-23

displaying device information, 22-15

displaying node information, 22-18

displaying OPCOM messages, 22-25

displaying status and counters, 22-23, 22-24

enabling MOP downline load service, 22-23

load trace facility, 22-25

MOP console carrier, 22-26

MOP downline load service management,
22-23

MOP trigger boot, 22-27

node database management, 22-17

running, 22-7

setting device information, 22-15

setting node information, 22-18

setting up LAN MOP, 22-22

SPAWN function, 22-9

using command files, 22-9

LANCP (Local Area Network Control Program)

See LAN Control Program utility

LAN drivers

address

multicast, 22-3

node, 22-3

physical, 22-3

addresses, 22-3

Ethernet, 22-2

FDDI, 22-2

port, 22-2

protocol type, 22-2

Token Ring, 22-2

Language formats, 5-42

Languages

specifying, 5-43, 5-44, 5-48

LANs (local area networks)

clearing counters, 22-25

connections, 21-8

deleting device information, 22-17

deleting node information, 22-20

device management, 22-9

disabling MOP downline load service, 22-23

LANs (local area networks) (cont'd)

displaying device information, 22-15

displaying LAN device configurations, 22-9

displaying LAN device parameters, 22-10

displaying node information, 22-18

displaying OPCOM messages, 22-25

displaying status and counters, 22-23, 22-24

enabling MOP downline load service, 22-23

LANACP device database management, 22-14

LANACP related OPCOM messages, 22-22

LAN Auxiliary Control Program (LANACP)

utility, 22-5

LAN Control Program (LANCP) utility, 22-6

LANCP command files, 22-9

LANCP SPAWN function, 22-9

LAN firmware updates, 22-14

LAN MOP and DECnet MOP, 22-20

load trace facility, 22-25

migrating DECnet MOP to LAN MOP, 22-20

MOP console carrier, 22-26

MOP downline load service, 22-20

MOP downline load service management,
22-23

MOP trigger boot, 22-27

node database management, 22-17

running the LANACP utility, 22-6

sample LAN MOP set up, 22-22

setting device information, 22-15

setting LAN device parameters, 22-12

setting node information, 22-18

setting up LAN MOP with CLUSTER_CONFIG,
22-21

stopping the LANACP utility, 22-6

system management enhancements, 22-1

LASTCP (LASTport Control Program)

See LASTport Control Program utility

LASTport Control Program (LASTCP), 23-5 to 23-13

account requirements, 23-12

command summary, 23-10

exiting, 23-9

functions, 23-12

Help facility, 23-10

invoking, 23-9

MAXBUF system parameter requirement,
23-12

privileges required, 23-9

LASTport Control Program utility

functions, 23-9

LASTport/Disk protocol, 23-5

LASTport/Disk service, 23-12

ESS\$DADDRIVER, 23-13

LASTport protocol, 23-5

LASTport/Tape protocol, 23-5

LASTport/Tape service, 23-12

ESS\$MADDRIVER, 23-13

- LASTport transport, 23-13
- LAT\$CONFIG.COM command procedure, 24-10
 - invoking during system startup, 5-15
- LAT\$STARTUP.COM command procedure, 24-10, 24-11
 - invoking during system startup, 5-15
- LAT\$SYSTARTUP.COM command procedure, 24-10, 24-11, 24-12, 24-17
 - example, 24-17
 - invoking during system startup, 5-15
- LATACP
 - See LAT Ancillary Control Process
- LAT Ancillary Control Process (LATACP), 24-19
- LAT Control Program (LATCP) utility, 24-3, 24-8
 - See also LAT software
 - DELETE QUEUE_ENTRY command, 24-16
 - exiting, 24-9
 - features, 24-9
 - invoking, 24-9
 - queuing incoming requests, 24-15
 - SET NODE command, 24-16
 - SET SERVICE command, 24-16
 - setting up limited services, 24-15
 - SHOW QUEUE_ENTRY command, 24-16
 - SHOW SERVICE command, 24-16
 - summary of commands, 24-9
- LATCP (Local Area Transport Control Program)
 - See LAT Control Program (LATCP) utility
- LAT server, 23-6
- LAT software
 - See also LAT Control Program utility
 - advantages and uses, 24-2
 - application programs, 24-2
 - creating a service, 24-13
 - customizing, 24-12
 - enabling outgoing connections, 24-16
 - load balancing, 24-3
 - managing the database size, 24-19
 - modems, 24-2
 - outgoing connections, 24-2, 24-3, 24-8
 - printers, 24-2
 - autostart queues on, 13-4, 13-10
 - increasing availability of, 13-4, 13-10
 - LATSYM symbiont, 13-3, 13-78
 - PRTSMB symbiont, 13-78
 - sample configuration, 13-10
 - setting up, 13-14
 - spooling, 7-14
 - troubleshooting, 13-78
- queuing incoming requests, 24-15
- service
 - announcements, 24-4, 24-8
 - database, 24-8
 - dedicated applications, 24-2
 - defined, 24-2
 - node, 24-3, 24-4, 24-9
 - remote printing, 24-2
- LAT software (cont'd)
 - setting up limited services, 24-15
 - setting up ports, 24-14
 - starting network in command procedure, 5-15, 24-11
 - starting with LAT\$STARTUP.COM, 5-15, 24-10, 24-11
 - terminals, 7-12, 24-2
 - determining characteristics of, 7-12
 - disconnecting, 7-11
- LATSYM symbiont, 13-3, 13-78
- Layered products
 - startup database, 5-18, 5-19
 - startup phases, 5-18
- Layered software installation, 3-18
- Level 1 routers
 - See Routers, Level 1
- Level 2 routers
 - See Routers, Level 2
- Lexical functions
 - F\$GETJPI, 26-9
 - F\$GETQUI, 13-51
 - F\$GETSYI, 26-9
 - getting information about queues, 13-51
 - getting information about vector processing, 26-9
- LIBDECOMP.COM command procedure, 16-6
- Libraries
 - See also Device control libraries
 - See Device control libraries
- License database, 3-8
 - logical name defining location, 5-8
- License Management Facility (LMF), 3-8, 5-5
- LICENSE MODIFY command, 3-8
- Licenses
 - loading, 3-8
 - in system startup, 5-5
- Limits
 - See Process limits
 - See UAFs (user authorization files), resource limits
- Line printer
 - UETP test image, 17-34
- Line printers
 - preparing for UETP, 17-3, 17-5, 17-8
 - testing with UETP, 17-31
 - UETP output, 17-33
 - UETP test image, 17-34
- Lines
 - communications
 - definition, 21-4
 - network, 21-5
 - definition, 21-4
 - overflow
 - controlling, 13-43

- Linkable images, 16-9
- LINK command
 - See also Linker utility
 - /SELECTIVE_SEARCH qualifier, 5-23
- Linker utility (LINK)
 - linking against SYS.STB, 5-23
- Links
 - logical
 - See Logical links
- List operations
 - with BACKUP, 10-19
- LMF\$JOURNAL file type, 25-3
- LMF\$LICENSE logical name
 - defining during system startup, 5-8
- Load balancing
 - using LAT software, 24-3
- LOAD command
 - in SYSGEN (VAX), 7-7
- Loading device drivers
 - automatically, 7-6, 7-7
 - manually, 7-8
 - on VAX, 7-7
- Load leveling
 - dynamic, 26-2
- LOADS logical name, 17-34
- Load tests
 - defining user load for UETP, 17-14
 - description, 17-34
 - error during UETP, 17-27
 - running individually, 17-13
- Load trace facility, 22-25
- Local Area Network Auxiliary Control Program
 - See LAN Auxiliary Control Program utility
- Local area networks
 - See LANs
- Local identifiers, 11-11
- Local nodes
 - definition, 2-13
- Local page and swap files
 - installing with SATELLITE_PAGE.COM
 - procedure, 5-6, 15-20
- Location
 - of page file, 15-3, 15-5
 - specifying alternate, 15-24
 - of queue database, 12-5, 12-7
 - master file, 12-6
 - queue and journal files, 12-6
 - changing, 12-6
 - of swap file, 15-5
 - specifying alternate, 15-24
 - of system dump file, 15-2
 - of system files
 - redefining with logical names, 16-8
- Location of software products, 3-28
- Log file generated by UETP
 - OLDUETP.LOG, 17-21
- Log files
 - generated by UETP
 - OLDUETP.LOG, 17-20
 - moving to reduce system disk I/O, 16-8
 - operator
 - creating new, 18-22
 - enabling and disabling classes, 18-23
 - maintaining, 18-24
 - printing, 18-24
 - restarting, 18-24
 - security alarm messages, 18-21
 - setting up, 18-22
 - specifying location, 18-22
 - troubleshooting the queue manager, 12-14
 - security audit, 18-2
 - creating new version, 18-30
 - reviewing, 18-26
- Log files generated by UETP
 - during the load test, 17-27
 - NETSERVER.LOG, 17-26
- Logging in, 2-9
 - See also Logins
 - to a remote host, 21-13
 - when errors in login procedures prevent, 4-8
 - when errors in startup procedures prevent, 4-8
 - when forgotten passwords prevent, 4-10
- Logging out
 - using command procedure, 6-22
 - while using VMSINSTAL.COM command
 - procedure, 3-7
- Logging startup
 - with SYSMAN, 4-15
- Logical links
 - network, 21-5
 - definition, 21-4
- Logical names
 - See also Symbols
 - ACCOUNTNG, 19-4
 - AGEN\$FEEDBACK_REQ_TIME, 14-21
 - assigning for shareable images, 16-17
 - assigning systemwide in system startup, 5-8
 - assigning to customize the SHUTDOWN.COM
 - procedure, 4-32
 - assigning to devices, 5-11
 - for software product installations, 3-28
 - for source and destination locations, 3-28
 - for system components
 - recommended privilege mode, 5-9
 - LMF\$LICENSE, 5-8
 - NETNODE_REMOTE, 5-8
 - NETPROXY, 5-8
 - OPC\$LOGFILE_CLASSES, 18-23
 - OPC\$LOGFILE_ENABLE, 18-23
 - OPC\$LOGFILE_NAME, 18-22, 18-23
 - OPC\$OPA0_ENABLE, 18-23
 - overriding source and destination locations,
 - 3-28

Logical names (cont'd)

- PCSI\$DESTINATION, 3-28
- PCSI\$SOURCE, 3-28
- privilege modes, 5-9, 16-15
- QMAN\$MASTER, 12-6
- redefining location of system files, 16-8
- RIGHTSLIST, 5-8
- SHOW_CLUSTER\$INIT, 20-14
- SHUTDOWN\$DECNET_MINUTES, 4-33
- SHUTDOWN\$DISABLE_AUTOSTART, 4-33, 13-56
- SHUTDOWN\$INFORM_NODES, 4-33
- SHUTDOWN\$MINIMUM_MINUTES, 4-33
- SHUTDOWN\$QUEUE_MINUTES, 4-33
- SHUTDOWN\$TIME, 4-33
- SHUTDOWN\$VERBOSE, 4-33
- specifying as mail addresses, 5-33
- STARTUP\$STARTUP_LAYERED, 5-18
- STARTUP\$STARTUP_VMS, 5-18
- SYS\$ANNOUNCE, 5-14
- SYS\$AUDIT_SERVER_INHIBIT, 5-9, 18-27
- SYS\$DECDTM_INHIBIT, 25-21
- SYS\$ERRORLOG, 5-8
- SYS\$JOURNAL, 25-3
- SYS\$MONITOR, 5-8
- SYS\$STARTUP, 5-2
- SYS\$SYLOGIN, 5-17
- SYS\$TIMEZONE_DIFFERENTIAL, 5-35
- SYS\$WELCOME, 5-15
- SYSUAF, 5-8
- translation of, 5-33
- trusted, 16-15
- UAFALTERNATE, 4-10
- using in SYSMAN, 2-14
- VMSMAIL_PROFILE, 5-8

Logical names used by UETP

- CTRLNAME, 17-32
- LOADS, 17-34
- SYS\$INPUT, 17-31
- SYS\$OUTPUT, 17-33

Logical name tables

- definition, 5-8

Logical queues

- assigning, 13-56
- description, 13-4
- recommended use, 13-4, 13-56

Login command procedures

- booting without, 4-8
- defining announcements in, 5-14
- defining location of during system startup, 5-17
- definition, 5-16, 5-17
- for captive account, 6-29
 - sample, 6-30
- for SYSTEM account, 5-16
- individual, 6-19, 6-20
 - sample, 6-21
- in SYSMAN, 2-16

Login command procedures (cont'd)

- LOGIN.COM, 5-16
- SYLOGIN.COM, 5-16
- systemwide, 5-16, 6-20
 - sample, 6-20
- to a remote host, 21-13
- user-specified, 6-20
- when errors prevent you from logging in, 4-8

Logins

- See also Logging in
- controlling number of dialup attempts, 11-6
- restricting by function, 6-28
- restricting by time, 6-27, 6-28
- sequence of events, 6-5

- Log in to a remote host, 21-13

- LOGOUT command, 6-23, 26-9

Logout command procedures

- SYLOGOUT.COM, 6-22

Long report format

- See Console report during UETP

Loopback tests

- network
 - circuit-level, 21-17
 - definition, 21-17
 - node-level, 21-17

Lost files

- recovering, 8-56 to 8-58

- LPBEGIN phase of system startup, 5-19

- LPBETA phase of system startup, 5-19

- LPMAIN phase of system startup, 5-19

- LTA_n devices, 7-11

- LTDRIVER (LAT port driver)

- turning on and off, 24-9

M

Machine check errors

- if returned when booting, 4-16

- MAD virtual tape unit, 23-12

Magnetic tape ancillary control process

- See MTACP

MAIL

- managing, 6-38

Mail utility (MAIL)

- logical name for, 5-8
- MAIL\$SYSTEM_FLAGS logical name, 5-32
- managing accounts, 6-38
- READ/NEW command, 5-33
- sending AUTOGEN report with, 14-22
- user profile entry
 - modifying, 6-38

Maintenance operation protocol

- See MOP

Management environment

- clusterwide, 2-15
- defining, 2-12 to 2-15
- individual nodes, 2-14

- Management environment (cont'd)
 - local and nonlocal environments, 2-13
- Manager, queue
 - See Queue managers
- Managing a multiprocessing environment, 26-3
 - tasks for, 26-3
- Managing a vector processing environment
 - tasks for, 26-5
- Managing devices
 - magnetic tape
 - tasks for, 7-16
 - printers
 - setting characteristics, 7-13
 - tasks for, 7-13
 - tasks for, 7-1
 - terminals
 - setting characteristics, 7-9
 - tasks for, 7-9
- Managing installed software, 3-34
- Managing page, swap, and dump files
 - tasks for, 15-1
- Managing performance, 16-1
 - See also Tuning
 - See Performance, managing
 - choosing a workload management strategy, 16-3
 - considering hardware capacity, 16-6
 - distributing work load, 16-4
 - evaluating tuning success, 16-6
 - installing images, 16-9
 - knowing your work load, 16-2
 - options for, 16-6
 - system tuning, 16-4
 - tasks for, 16-1
 - with vector processing, 26-7
- Managing system parameters
 - tasks for, 14-1
- Managing the LAT database size, 24-19
- Managing the queue manager and queue database, 12-1
- Mandatory updates
 - definition, 3-5
- Mapping pointers
 - resetting for windows, 8-24
- Marginal vector consumer
 - detection of, 26-8
- Margin size
 - specifying in forms, 13-42
- Mass storage control protocol
 - See MSCP
- Master command procedure
 - See UETP.COM procedure
- Master file directories
 - See MFDs
- Master file of queue database, 12-4
 - See also Queue database
 - location
 - specifying, 12-6
 - mounting of disk holding, 12-6
 - QMAN\$MASTER logical name, 12-6
 - saving, 12-12
- MAXACCTJOBS process limit, 6-43
- MAXDETACH process limit, 6-43
- Maximum account jobs process limit, 6-43
- Maximum detached process limit, 6-43
- MAXJOBS process limit, 6-43
- MAXSYSGROUP system parameter, 11-8
- MAX_DUMPFILE symbol, 15-24
- MAX_PAGEFILE_n_SIZE symbol, 15-24
- MAX_PAGEFILE symbol, 15-24
- MAX_ prefix for AUTOGEN, 14-20
- MAX_SWAPFILE_n_SIZE symbol, 15-24
- MAX_SWAPFILE symbol, 15-24
- Media errors
 - analyzing, 8-64
- Memory
 - allotted to vector consumer processes, 26-7
 - conserving with shareable images, 16-11
 - images in, 16-11
 - information captured in crash dump, 15-2
 - physical dump, 15-3, 15-10
 - selective dump, 15-3, 15-10
 - making efficient use of by installing images, 16-8
 - paging, 15-4
 - sections in, 16-12
 - swapping, 15-4
 - when large sizes prevent storing a complete system dump, 15-10
- Messages
 - broadcast
 - removing, 20-10
 - defining welcome, 5-14
 - DIBOL
 - starting the DIBOL message manager, 5-16
 - enabling and disabling, 18-17
 - error
 - removing, 20-10
 - error log, 18-6
 - indicating execution of startup command procedures
 - site-independent, 4-5
 - site-specific, 4-5
 - indicating high page or swap file fragmentation, 15-27
 - indicating insufficient page file size, 15-9
 - indicating lack of installed page file, 5-6
 - indicating successful boot, 4-5
 - indicating that a vector processor is not available, 26-6
 - indicating that login is possible, 4-5

Messages (cont'd)

- login welcome, 2-10

OPCOM

- See OPCOM messages

- operator replies, 2-25, 18-19

- operator requests, 2-25

- question mark (?), 4-16

- saving in a file, 3-29

- security alarm, 18-21

- sending to users with OPCOM, 2-22

- suppressing the display, 14-12

- user requests, 18-19

- using when managing a system, 2-5

- WRITEBOOT, 4-19

MFD (master file directory)

- BACKUP stores save-set file in, A-8

- contains directory structure for volume set, 8-33

- definition, A-8

- on root volume in volume set, 8-31

- reserved file, A-8

- reserved files listed in, 8-4

Minimerge and volume shadowing, 10-42

Minimum startup

- booting with, 4-14

- MIN_DUMPFILEx symbol, 15-24

- MIN_PAGEFILEn_SIZE symbol, 15-24

- MIN_PAGEFILE symbol, 15-24

- MIN_prefix for AUTOGEN, 14-20

- MIN_SWAPFILEn_SIZE symbol, 15-24

- MIN_SWAPFILE symbol, 15-24

- MODE logical name, 17-18, 17-38

- Mode menu, 3-37

Modes

- See Executive mode

- See Privilege mode

Modifiable startup command procedure

- See Site-specific startup command procedure

Modifying size of page, swap, and dump files

- See also Changing size of page, swap, and dump files

- See Changing size of page, swap, and dump files

Modifying system parameters

- See also Changing system parameters

- See Changing system parameters

MODPARAMS.DAT file, 14-17, 14-18

- ADD_prefix, 14-19

- controlling page, swap, and dump files, 15-18, 15-22

- specifying location, 15-24

- specifying size of individual files, 15-24

- specifying size of total file space, 15-23

- controlling parameter values set by AUTOGEN, 14-4, 14-18

- creating page, swap, and dump files, 15-17, 15-24

MODPARAMS.DAT file (cont'd)

- including external parameter files in, 14-22

- increasing parameter values, 14-19

- MAX_prefix, 14-20

- MIN_prefix, 14-20

- sample, 14-17

- specifying an alternate default startup command, 4-13

- specifying parameter values

- absolute, 14-20

- maximum, 14-20

- minimum, 14-20

- storing your system parameter changes in, 14-5

Modules

- device control

- See Device control modules

MONITOR.COM command procedure, 18-40

- used with Monitor utility, 18-39

MONITOR command

- See also Monitor utility

- controlling amount of time between displays of screen images, 18-37

- recording data on time spent, 18-36

- recording file system and process data, 18-36

- routing display output, 18-35

- specifying how log request is to run, 18-34

- specifying input file, 18-36

- specifying node, 18-35

- specifying time of display, 18-36

- storing display output, 18-37

Monitoring a multiprocessing environment, 26-3

Monitor utility (MONITOR)

- See also MONITOR command

- analyzing disk use with, 8-8

- class types, 18-31

- command procedures, 18-39

- for cluster summaries, 18-40

- to initiate continuous recording, 18-41

- to produce cluster summaries, 18-42

- description, 18-31

- directing display, 18-33

- displaying and recording concurrently, 18-36

- displaying live monitoring, 18-34

- displaying network information, 21-17

- entering commands, 18-34

- exiting from, 18-34

- invoking, 18-33

- logical name for, 5-8

MONITOR.COM command procedure

- using to create summary file, 18-39

MONSUM.COM command procedure

- using to generate clusterwide multifile

- summary reports, 18-39

- moving log file to reduce system disk I/O, 16-8

- parameters, 18-31

- playing back monitoring, 18-37

- playing back remote monitoring, 18-38

Monitor utility (MONITOR) (cont'd)

- producing reports, 18-39
- qualifiers, 18-33
- recording live monitoring, 18-35
- reports, 18-40
- rerecording monitoring, 18-39
- running continuously, 18-39
- SUBMON.COM procedure
 - using to start MONITOR.COM as a detached process, 18-39
- version compatibility, 18-43
- MONSUM.COM command procedure, 18-42
 - used with Monitor utility, 18-39
 - using to generate clusterwide multifile summary reports, 18-39
- MOP downline load service
 - clearing counters, 22-25
 - console carrier, 22-26
 - disabling, 22-23
 - displaying status and counters, 22-23, 22-24
 - enabling, 22-23
 - LANACP server, 22-23
 - LAN MOP, 22-20
 - coexisting with DECnet MOP, 22-20
 - migrating DECnet MOP to LAN MOP, 22-20
 - sample setup, 22-22
 - setup with CLUSTER_CONFIG.COM, 22-21
 - trigger boot, 22-27
- MOP protocol
 - downline loading, 23-2

Motif

See DECwindows Motif interface

MOUNT command

See also Mounting volumes

- adding to volume sets, 8-22
- assigning a volume set name, 8-31
- avoiding use of /CLUSTER with SYSMAN DO command, 20-21
- controlling whether header labels are written to a volume, 8-24
- creating a public volume, 8-24
- creating disk volume sets, 8-30
- creating volume sets, 8-22
- disabling automatic notification of mount failures, 8-22
- disabling automatic volume switching, 8-40
- disabling mount verification feature for disks, 8-23
- disabling mount verification feature for tapes, 8-24
- disk and tape volumes, 8-20
- enabling automatic notification of mount failures, 8-22
- enabling mount verification feature for disks, 8-23
- enabling mount verification feature for tapes, 8-24

MOUNT command (cont'd)

- enabling the processing of subsystem ACEs, 8-23
- enabling the write cache for a tape device, 8-24
- ensuring that tape volume set has been initialized, 8-40
- establishing default file attributes for records on ISO 9660 media, 8-24
- foreign volumes, 8-22
- for foreign volumes, 8-24, 9-13
- including a quoted text string as part of mount request, 8-22
- inhibiting access checks, 8-25
- in system startup
 - for remote InfoServer devices, 23-14
 - mounting page and swap file disks, 5-6
 - special consideration about operator assistance, 5-11
- in VMScluster environment, 8-21, 8-22
- ISO 9660 media, 8-23
- overriding expiration date, 9-14
- overriding protection, 8-25
- overriding protection checks, 8-23
- overriding the volume identification field, 8-25
- overriding UIC in second volume label, 8-25
- parameters, 8-21
- protection codes, 9-12
- public volumes, 8-21
- qualifiers, 8-22, 8-24
- requesting operator assistance, 8-22
- resetting the number mapping pointers, 8-24
- specifying block size for tape, 8-24
- specifying number of bytes in each record, 8-26
- specifying record size, 8-26
- specifying that other users can access current volume, 8-23
- specifying the number of directories that the system keeps in memory, 8-22
- specifying the number of disk blocks allocated, 8-22
- specifying UIC, 8-25
- suspending quota operation on a volume, 8-52
- tape, 9-21
- to mount disk holding page and swap files, 15-20

Mounted forms

matching stock, 13-42

Mounting forms, 13-64

Mounting of queue database disk, 12-6

Mounting volumes, 10-15

See also MOUNT command

disks, 8-20

for queue database files, 12-6

in system startup, 5-11

early, 5-12

for page and swap files, 5-6

InfoServer, 23-14

- Mounting volumes
 - disks
 - in system startup (cont'd)
 - special consideration about operator assistance, 5-11
 - virtual device unit, 23-14
 - if device is unavailable, 8-27
 - in a VMScluster environment, 8-21
 - operator assistance, 8-18
 - public, 8-21
 - substituting, 8-27
 - tape, 8-20
 - tape volume sets, 8-36
 - Mounting volume sets
 - See also MOUNT command
 - disk, 8-31, 8-32, 8-33
 - tape
 - with automatic volume switching disabled, 8-40
 - Mount messages
 - disabling with SUBSYSTEM qualifier, 8-23
 - Mount utility (MOUNT)
 - using to mount ISO 9660 volume sets, 8-34
 - Mount verification, 8-58
 - aborted
 - OPCOM message, 8-63
 - aborting by dismounting, 8-62
 - canceling, 8-62, 8-63
 - definition, 8-58
 - device off line, 8-59, 8-61
 - device write-locked, 8-61
 - enabling, 8-60
 - enabling and disabling, 8-23
 - for tapes, 8-24
 - messages, 8-60
 - operation of, 8-59
 - timeout, 8-60
 - OPCOM message, 8-60
 - MOVE keypad function, 20-7
 - Moving jobs from one queue to another, 13-57
 - MSCP (mass storage control protocol), 20-2
 - MSGHLP
 - See Help Message utility
 - .MSGHLP\$DATA files
 - adding to the Help Message database, 5-28
 - MSGHLP\$LIBRARY.MSGHLP\$DATA file, 5-28
 - MSGHLP\$LIBRARY logical name, 5-28
 - MTACP (magnetic tape ancillary control process)
 - definition, 8-6
 - Multiple queue managers, 12-3, 12-10
 - commands affected by, 12-4
 - managing queues with, 12-10
 - moving queues, 12-10
 - restriction, 12-3
 - specifying names of, 12-4
 - specifying queue manager name, 12-4, 12-10
 - use of queue database, 12-4

- Multiple-site network, 21-9
- Multiprocessing, 26-2
 - definition, 26-2
 - displaying information, 26-3
 - monitoring, 26-3
 - system parameters, 26-3
- MULTIPROCESSING system parameter, 26-3
- MVTIMEOUT system parameter, 8-60, 8-62

N

- Namespaces
 - in DECnet/OSI full name, 21-3
- Naming conventions
 - for queue and journal files for additional queue managers, 12-5
 - of devices, 7-1
 - in a VMScluster environment, 7-1
 - virtual terminals, 7-11
- NCP (Network Control Program)
 - commands
 - CLEAR, 21-11
 - DEFINE, 21-11
 - LIST, 21-15
 - PURGE, 21-11
 - SET, 21-11
 - SHOW, 21-15
 - controlling proxy access, 6-37
 - display commands, 21-15
 - Ethernet configurator, 21-17
 - testing network, 21-17
 - using to build or modify configuration database, 21-11
- NET\$PROXY.DAT file, 11-7
- NET\$PROXY.DAT files, 6-4, 6-34
 - moving to reduce system disk I/O, 16-8
- NET\$PROXY logical name
 - defining to reduce system disk I/O, 16-8
- NETCONFIG.COM command procedure
 - configuring a node automatically, 21-11, 21-12
- NETDRIVER (network driver)
 - connecting, 7-7, 7-8
- NETNODE_REMOTE logical name
 - defining during system startup, 5-8
- NETPROXY.DAT files, 6-4, 6-34
 - moving to reduce system disk I/O, 16-8
- NETPROXY logical name
 - defining during system startup, 5-8
 - defining to reduce system disk I/O, 16-8
- Network communications device
 - connecting, 7-7, 7-8
- Network Control Program
 - See NCP
- Network identifiers, 11-11
- Network interface
 - See DECnet

Network proxy accounts, 6-34

Networks

See also DECnet

See also Nodes

area routing, 21-6

becoming node in, 21-11

bridge, 21-7

channel, 21-5

circuit

definition, 21-4

communications device, 7-8

communications line, 21-12

configuration command procedure

See NETCONFIG.COM

configurations, 21-6

See also Configurations, network

connecting to existing, 21-12

connections, 21-8

local area, 21-8

multipoint, 21-9

point-to-point, 21-8

wide area, 21-8

worldwide, 21-9

definition, 21-4

end node, 21-6

Ethernet

definition, 21-5

getting information about nodes, 21-10

interface

OpenVMS, 21-7

Internet, 21-13

LAN enhancements, 22-1

line

definition, 21-4

logical link

definition, 21-4

loopback tests, 21-17

managing a node, 21-14

monitoring tools, 21-15

DTS/DTR, 21-17

Monitor utility, 21-17

NCP display commands, 21-15

NCP Ethernet configurator, 21-17

multinode, 21-5

multiple-area, 21-6

multiple-site, 21-9

nodes

configuring, 21-12

connecting, 21-8

definition, 21-4

object

definition, 21-4

planning configuration, 21-12

proxy database

creating, 6-35

logical name defining location, 5-8

remote node database

Networks

remote node database (cont'd)

logical name defining location, 5-8

router, 21-5

routing

adaptive, 21-5

definition, 21-5

save sets, 10-7

security

providing for, 21-12

starting in system startup, 5-15

testing

using NCP (Network Control Program),

21-17

verifying connection to, 21-12

Node names

assigning, 21-3

DECnet/OSI full names, 21-2

DECnet for OpenVMS, 21-2

required for OpenVMS InfoServer Client

startup, 23-10

setting up a system for assigning, 21-3

specifying in MAIL, 5-33

Nodes

See also Networks

adjacent, 21-4

configuring in a network, 21-12

connecting to network, 21-8

definition, 21-4

deleting node information, 22-20

end, 21-6

establishing in network, 21-11

getting information about other nodes in

network, 21-10

in LAT database, 24-8

monitoring, 21-14

multiple network, 21-5

nonrouting, 21-6

preparing system to become network node,

21-12

providing security for, 21-12

router definition, 21-5

routing, 21-6

setting node information, 22-18

tools to monitor network, 21-15

transferring files between, 9-24

Noneductible resource, 6-3

Nonrouting nodes

See End nodes

Nonstop boot

definition, 4-3

performing, 4-3

NPAGEDYN system parameter, 26-7

NUM_ETHERADAPT symbol, 14-21

NUM_NODES symbol, 14-21

O

Objects

network

definition, 21-4

protecting volume, 8-14

OLDUETP.LOG file, 17-20

OPA0: device, 2-23

OPC\$LOGFILE_CLASSES logical name, 18-23

OPC\$LOGFILE_ENABLE logical name, 18-23

OPC\$LOGFILE_NAME logical name, 18-23

for operator log file, 18-22

OPC\$OPA0_ENABL logical name, 18-23

OPCCRASH.EXE program, 4-27

OPCOM

messages

displayed during LANACP start, 22-22

generated by LANACP, 22-25

OPCOM (Operator Communication Manager), 10-17

See also Operator log file

automatic restart, 2-22

classes

enabling, 2-23

enabling and disabling for log file, 18-23

communicating with operators, 8-18

communicating with users, 8-18

components of, 2-19

default behavior, 2-21

disabling operator terminals, 2-24

enabling operator classes, 2-23

enabling operator terminals, 2-23

failure, 2-22

illustration, 2-19

log file, 2-21, 18-17

messages

See OPCOM messages

mount verification, 8-60

operator log file

See Operator log file

operator terminals, 2-21

printing operator log file, 18-24

process, 2-21

creation during system startup, 5-5

process dump file, 2-22

replying to operator requests, 2-25

requirements, 2-21

restarting operator log file, 18-24

sending requests to an operator, 2-25

setting up operator log file, 18-22

specifying default state of operator log file, 18-23

startup of, 2-22

uses of, 2-19

using to request assistance, 10-17

OPCOM.DMP process dump file, 2-22

OPCOM messages, 2-21

continuation volume request, 8-40, 9-24

controlling, 2-22

mount request message, 8-26

mount verification, 8-60

aborted message, 8-63

timeout message, 8-60

operator reply, 18-19

security alarm, 18-21

sending, 2-22

SYSGEN, 18-20

user request, 18-19

Open file limit, 6-43

Open images, 16-10, 16-12

OpenVMS Management Station

description, 2-2

documentation, 2-3

features, 2-4

Operating systems

building on another disk, 2-27

copying to another system disk, 2-29

installing, 3-4

updating, 3-5

upgrading, 3-5

Operation button on POLYCENTER Software

Installation utility, 3-37, 3-39

OPERATOR.LOG file

See Operator log file

Operator assistance

operator classes, 2-23

replying to operator requests, 2-25

with MOUNT command, 8-18

Operator classes

See OPCOM, classes

Operator communication manager

See OPCOM

Operator consoles

enabling in system startup, 5-5

Operator log file

See also OPCOM

See also OPCOM messages

closing current and opening new, 18-17

creating new, 18-22

definition, 2-21

device status message, 18-17, 18-22

enabling and disabling classes, 18-23

enabling in system startup, 5-5

initialization message, 18-17

logging user requests in, 18-19

maintaining, 18-24

printing, 18-22, 18-24

purging, 18-24

during system startup, 5-14

recording changes to system parameters, 18-20

request identification number, 18-19

response recorded in, 18-19

Operator log file (cont'd)

- restarting, 18-24
- security alarm messages in, 18-21
- setting up, 18-16, 18-22
- specifying default state, 18-23
- terminal enable and disable message, 18-22
- troubleshooting the queue manager, 12-14
- user request and operator reply messages, 18-22
- volume mount and dismount messages, 18-22

Operators (computer)

- assistance with MOUNT command in system startup, 5-11
- classes of, 2-23
- replying to requests, 2-25
- requesting assistance from, 8-18
- sending requests to, 2-25

Operator terminals, 2-21

- designating, 2-23
 - in batch or SYSTARTUP, 2-21
- enabling and disabling, 2-23, 2-24, 18-17
 - security alarms, 18-29
- security terminal, 18-21
- setting up, 8-18
- user request, 8-18

Optional files

- adding and removing, 5-1

Option list

- parameter for VMSINSTAL.COM command procedure, 3-12

Options

- for software product installations, 3-30

Outgoing LAT connections, 24-3, 24-8

- enabling with LAT software, 24-16

Output a newly created PCF, 3-30

Output devices

- See Printers
- See Terminals

Output during UETP

- terminal and line printer, 17-33

Output execution queues

- See also Execution queues
- See also Output queues
- definition, 13-3

Output jobs

- See also Output queues
- aligning forms, 13-77
- allowing to complete before stopping a queue, 13-55
- changing scheduling priority, 13-72
- controlling, 13-68
- controlling print position and alignment, 13-75, 13-76, 13-77
- deleting, 13-74
- holding and releasing, 13-71
- modifying, 13-70
- monitoring, 13-68

Output jobs (cont'd)

- requeuing
 - executing, 13-73
 - pending, 13-73
- resuming printing, 13-75, 13-76, 13-77
- retaining in a queue, 13-74
- scheduling, 13-32, 13-72
- status

See Job status

- suspending, 13-75

Output queues

See also Output jobs

See also Output queuing environment
allowing jobs to complete before stopping, 13-55

- assigning a default form, 13-63
- canceling characteristics, 13-59
- changing DEFAULT form, 13-63
- commands for managing, 13-46
- controlling line overflow in forms, 13-42
- creating, 13-15
- defining a form, 13-61
- deleting, 13-57
- execution

- description, 13-3
- printer, 13-3
- server, 13-3
- terminal, 13-3

mounting a form on, 13-64

on standalone workstations, 13-8

options, 13-18

- banner pages, 13-33
- characteristics, 13-28
- controlling page and line overflow, 13-43
- device control libraries, 13-44
- forms, 13-42
- qualifiers for specifying, 13-20 to 13-21
- restricting access, 13-22
- retaining jobs, 13-26

order of device control module output, 13-45

pausing, 13-53, 13-75

- to align position of print for preprinted forms, 13-75, 13-77

- to change position of print, 13-75

rerouting jobs in, 13-56

specifying page and margin size in forms, 13-42

starting, 13-15

status, 13-50

stopping, 13-54, 13-55

- before shutting down a node, 13-55

Output queuing environment

- for LAT printers, 13-10
- for mixed printers, 13-9
- for multiple printers of the same kind, 13-11
- in VMScluster environments, 13-12
- on a standalone workstation, 13-8

- Output queuing environment (cont'd)
 - sample configurations, 13-8 to 13-14
 - single printer, 13-8
 - spooled printers, 13-13
 - steps for setting up, 13-14

- Output specifier
 - to the BACKUP command, 10-4

- Overdraft limit
 - user exceeding quota, 8-49

- Overflow
 - See Lines, overflow
 - See Page overflow

- Overriding DECnet parameters, 14-21

- Owner
 - category of user access, 11-8

- Ownership
 - file
 - displaying, 9-7

P

- PAGEFILE.SYS file, 15-2, 15-4

- See also Page files
 - as system dump file, 15-2, 15-7
 - required location, 15-3
 - requirement
 - location, 16-8

- PAGEFILE n _NAME symbol, 15-17, 15-24

- PAGEFILE n _SIZE symbol, 15-18, 15-24

- Page files

- as system dump file, 15-2, 15-7, 15-17
 - releasing dump from, 15-16
 - size required for, 15-4

- changing sizes
 - with SWAPFILES.COM, 15-25

- creating
 - with AUTOGEN, 15-17, 15-24
 - with SYSGEN, 15-18

- definition, 15-4

- deleting after creating a new version, 15-28

- displaying, 15-6
 - size calculated by AUTOGEN, 15-17, 15-22

- fragmentation of, 15-27

- freeing dump information from, 15-4, 15-16

- installing
 - in system startup, 5-6, 15-5, 15-19, 15-21
 - when resized with AUTOGEN, 15-18, 15-22
 - with SYPAGSWPFILES.COM procedure, 15-20
 - with SYSGEN, 15-19

- limiting usage of, 15-9

- location
 - specifying for individual files, 15-24
 - message
 - indicating high fragmentation, 15-27
 - indicating insufficient size, 15-9

- Page files

- message (cont'd)
 - indicating lack of installed, 5-6
 - monitoring usage of, 15-6
 - mounting disk during system startup, 5-6, 15-20

- moving to improve performance, 16-8

- on a satellite, 15-20

- primary, 15-5

- purging, 15-28

- releasing dump from, 15-16

- requirements
 - location, 15-3, 15-5, 16-8
 - size for saving dumps, 15-4

- saving dump contents on reboot, 15-2

- secondary, 15-5, 15-20

- sizes

- See Page file sizes

- tasks for managing, 15-1

- VMScluster satellite node, 5-6

- writing crash dumps to, 15-2

- Page file sizes

- calculating

- for paging, 15-8

- for saving dumps, 15-7

- manually, 15-7, 15-8

- with AUTOGEN, 15-5

- changing

- recommended method, 15-22

- when to increase size, 15-9

- with AUTOGEN, 15-22

- with SYSGEN, 15-26

- determining current, 15-23

- displaying AUTOGEN's calculations, 15-17, 15-22

- message indicating insufficient, 15-9

- required

- for paging, 15-8

- for saving dumps, 15-4, 15-7

- specifying

- for individual files, 15-22, 15-24

- total for multiple files, 15-22, 15-23

- when to increase, 15-9

- Pagelets

- size, 6-39

- size of, 6-39

- Page overflow

- controlling, 13-44

- Pages

- size, 6-39

- Page setup modules, 13-44

- See also Device control modules

- specifying forms, 13-42

- Page width and length

- specifying in forms, 13-42

- Paging, 15-4
 - increased with vector processing, 26-7
- Paging file process limit, 6-43
- PAKs (Product Authorization Keys)
 - installing before using VMSINSTAL.COM
 - command procedure, 3-8
 - loading in system startup, 5-5
 - preventing nodes from sharing, 3-8
 - registering for DECnet, 21-12
- PAN command, 20-6
- PAN keypad function, 20-7
- Paper
 - See Printer paper
 - See Stock
- Paper jam
 - pausing printer to fix, 13-75
- Parameter files
 - See also System parameters
 - ALPHAVMSSYS.PAR (Alpha), 14-3
 - initializing parameters at boot time, 14-36
 - role in the boot process, 4-2
 - booting with alternate, 4-7
 - changing
 - effects of, 14-33
 - with SYSGEN, 14-33
 - with SYSMAN, 14-28
 - common in a VMScluster environment, 14-22
 - creating new
 - with SYSGEN, 14-34
 - default, 4-6, 14-3
 - including in MODPARAMS.DAT file, 14-22
 - limitation on error checking in, 14-19
 - MODPARAMS.DAT, 14-17
 - sample, 14-17
 - multiple, with AUTOGEN, 14-22
 - VAXVMSSYS.PAR (VAX), 14-3
 - initializing active parameters at boot time, 14-36
 - role in the boot process, 4-2
- Parameters
 - passing to a startup command procedure, 5-19
 - system
 - See also System parameters
 - See System parameters
- PARAMETERS command, 2-11
 - See also System parameters
 - in SYSMAN, 14-25
 - summary, 14-30
- Partition, 23-2
- \$PASSWORD card, 7-17
- Password generators
 - using to obtain initial password, 11-2
- Password management
 - forced change, 11-4
 - guidelines for protecting, 11-5
 - how to preexpire, 11-3
 - length of passwords, 11-4
- Password management (cont'd)
 - reasons to assign system passwords, 11-3
 - secondary passwords, 11-4
 - setting the expiration for, 11-4
 - setting up initial passwords, 11-2
 - when to use dual passwords, 11-2
- Passwords
 - See also Password management
 - conditions required for SYSMAN, 2-13
 - entering when logging in, 2-10
 - forgotten by user, 6-23
 - for VMScluster access
 - modifying, 20-16
 - history list, 11-2
 - InfoServer system, 23-6
 - modifying system, 6-9
 - modifying user, 6-13
 - when forgotten prevents you from logging in, 4-10
 - with SYSMAN, 2-16
- PATHWORKS
 - startup restrictions, 23-12
- Paused queue status, 13-51
- Pausing an output queue, 13-75
 - to align position of print for preprinted form, 13-75, 13-77
 - to change position of print, 13-75
- Pausing queue status, 13-51
- PCF (product configuration file), 3-20
 - command to create, 3-23
 - command to input, 3-30
 - command to output, 3-30
 - creating before installation, 3-22
 - creating during installation, 3-30
 - naming, 3-23
 - use existing, 3-30
- PCSI
 - See POLYCENTER Software Installation utility
- PCSI\$CONFIGURATION, 3-23
- PCSI\$DESTINATION, 3-28
- PCSI\$SOURCE, 3-28
- PDF (product description file), 3-20
- Pending bad block log file
 - BADLOG.SYS, A-9
 - reserved file, A-9
- Pending jobs
 - requeuing, 13-73
 - troubleshooting, 13-79
- Pending job status
 - definition, 13-69
 - deleting a job in, 13-74
 - determining whether a job is in, 13-79
 - inducing with STOP/QUEUE/REQUEUE
 - command, 13-73
 - requeuing a job in, 13-73
- Percent sign (%)
 - as wildcard character
 - using with tape volumes, 9-16

- Performance
 - See also Managing performance
 - See also Tuning
 - disk, 8-28
 - effect of file extension on, 16-7
 - importance of correct page file size, 15-8
 - importance of correct swap file size, 15-9
 - importance of sufficient hardware capacity, 16-6
 - improving
 - decompressing system libraries, 16-6
 - designing efficient applications, 16-4
 - disabling high-water marking, 16-7
 - encouraging batch processing, 16-4
 - for vector processing with batch queues, 26-7
 - installing frequently used images, 16-7, 16-9
 - moving page and swap files off system disk, 15-5
 - options for, 16-6
 - reducing system disk I/O, 16-8
 - relinking images, 16-7
 - restricting the number of interactive users, 16-4
 - restricting user login hours, 16-4
 - setting RMS file extend parameters, 16-7
 - slicing shareable images, 16-12
 - tuning the system, 16-4
 - with AUTOGEN feedback, 14-10
 - load balancing on public volumes, 8-9
 - monitoring, 16-2
 - tools used for, 16-2
 - of vector processing, 26-4
 - testing for disks, 8-8
- Performance management
 - using MONITOR, 18-39
- Permanent databases
 - network, 21-11
- PGFLQUOTA process limit, 6-43
 - limiting page file usage with, 15-9
 - minimum recommended value, 15-9
 - value for efficient backups, 10-10
- Phase controller for UETP
 - See UETPHAS00.EXE file
- Physical dump, 15-3
 - compared to selective dump, 15-3, 15-10
- POLYCENTER Software Installation utility, 3-18
 - See also PRODUCT command
 - commands and operations, 3-19
 - creating a PCF, 3-22
 - databases, 3-20
 - DCL interface, 3-18
 - DECwindows Motif interface, 3-18
 - exiting from DECwindows Motif, 3-35
 - getting help, 3-35
 - installing layered products, 3-6
 - POLYCENTER Software Installation utility (cont'd)
 - installing software, 3-27
 - installing the OpenVMS Alpha operating system, 3-2
 - interfaces to use, 3-18
 - managing installed software, 3-34
 - preliminary steps, 3-28
 - PRODUCT command to run, 3-18
 - product configuration files, 3-20
 - registering products, 3-27
 - registering volume labels, 3-33
 - removing installed products, 3-34
 - starting the utility, 3-18
 - upgrading the OpenVMS Alpha operating system, 3-2
 - using a product database, 3-26
 - using to add or delete optional system files, 5-2
- Pooled resource, 6-3
- Ports
 - setting up LAT, 13-14
- Position of print
 - aligning for preprinted forms, 13-75, 13-77
 - changing, 13-75, 13-76
- PostScript printing, 13-9
- PRCLM process limit, 6-43
- Preventing autostart queues from starting, 13-54
- Primary bootstrap image
 - definition, 4-17
 - role in boot process, 4-2
- PRIMARY day
 - defining for accounts, 6-27
- Primary page file, 15-5
 - See also PAGEFILE.SYS file
 - See also Page files
 - location requirement, 16-8
- Primary processors, 26-2
- Primary swap file, 15-5
 - See also SWAPFILE.SYS file
 - See also Swap files
- PRINT command
 - bypassing symbiont formatting, 13-44
 - overriding default form-feed options with, 13-44
 - preventing users from executing, 13-53
 - processing of, 13-2
 - specifying a form, 13-42
 - specifying banner pages, 13-60
 - specifying characteristics, 13-28
 - specifying job retention, 13-74
 - specifying scheduling priority, 13-72
 - specifying setup and page setup modules, 13-67
- Printer paper
 - controlling with forms, 13-42, 13-64
 - pausing to align, 13-75
 - sheet feed, 13-62

Printer paper (cont'd)

- specifying stock, 13-62
- specifying width, 13-61

Printers

- controlling functions of, 13-44

LAT

- See LAT software, printers

managing

- tasks for, 7-13

- setting characteristics, 7-13, 13-14

- in system startup, 5-12, 7-13

- setting up before creating queues, 13-14

spooled, 13-14

- definition, 7-14

- despooling, 7-15

- recommended use, 7-14

- requirement, 7-13

- sample configuration, 13-13

- spooling, 7-14

- testing spooling of, 7-16

- troubleshooting, 13-78

Print forms

- See Forms

Printing

- distributed, 13-13

- from applications using spooled printers, 7-14

- job status, 13-69

- PostScript, 13-9

- remotely, 13-13

- resuming at a specified position, 13-75, 13-76, 13-77

Print jobs

- See also Batch jobs

- See also Output jobs

- See also Print queues

- See Output jobs

Print queues

- See also Batch queues

- See also Output queues

- See Output queues

Print symbionts

- See Symbionts

Priority, 6-2, 6-31

- base, 6-2, 6-31

- choosing for a batch queue, 13-30

- effect of changing, 13-6

- specifying for a batch queue, 13-29

- specifying for a queue, 13-19

- job scheduling, 13-72

- See also Job scheduling

- changing for a job, 13-72, 13-73

- display on banner pages, 13-35, 13-37, 13-40, 13-41

- specifying for a job, 13-70, 13-72

Private volumes, 8-9

Privileged images, 16-10, 16-14

Privilege mode

- recommended for logical names of system components, 5-9

Privileges

- all, 6-4

- allowing nonprivileged users to run programs that require, 16-14

- changing in SYSMAN, 2-17

- enhancement for installed files, 16-13
- files, 6-4

- for SYSTEM account, 2-9

- GRPPRV, 11-8

- process, 6-3

- required for UETP, 17-23

- required to mount volumes, 8-15

SECURITY

- to mount volume with protected subsystems, 8-15

- SYSNAM, 8-16

- SYSPRV, 8-16

- giving rights of system user, 11-8

- VOLPRO, 8-14

- to mount volume as foreign, 8-15

Problems

- See also Troubleshooting

booting

- fixing by booting with default parameter values, 4-8

- fixing by booting with minimum startup, 4-14

- hardware, 4-16

- invalid boot block, 4-17

- devices not recognized by the system, 7-7

- forgotten password

- fixing by booting without the UAF, 4-10

- logging in, 4-8, 4-10

- OPCOM failure, 2-22

Processes

- maintaining when disconnecting a terminal, 7-11

- priority, 6-31

- See also Priority, base

- See also Priority, job scheduling

Processing environments

- multiprocessing

- See Multiprocessing

- vector processing

- See Vector processing

Processing job status, 13-69

Process limits, 6-2

- account jobs, 6-43

- adjusting for vector processing, 26-7

- AST queue, 6-41

- CPU default

- Process limits
 - CPU default (cont'd)
 - specifying a value for batch queues, 13-19, 13-29, 13-31
 - CPU maximum
 - specifying a value for batch queues, 13-19, 13-29, 13-31
 - CPU time, 6-42
 - detached process, 6-43
 - direct I/O count, 6-42
 - enqueue quota, 6-42
 - expiration, 6-42
 - for Snapshot facility, 4-20
 - jobwide logical name table, 6-43
 - open file, 6-43
 - paging file, 6-43
 - process jobs, 6-43
 - setting, 6-40
 - subprocess creation, 6-43
 - system resources, 6-39
 - timer queue entry, 6-43
 - working set
 - quota, 6-44
- Processors
 - adding to a multiprocessing active set, 26-3
 - removing from a multiprocessing active set, 26-3
- Process quotas
 - see Process limits
 - recommended values for backups, 10-10
 - setting before a backup, 10-9
 - working set
 - extent, 6-44
- Product Authorization Keys
 - See PAKs
- PRODUCT command
 - See POLYCENTER Software Installation Utility
 - DCL interface syntax, 3-18
 - DECwindows Motif interface syntax, 3-18
- Product configuration file
 - See PCF
- Product database, 3-35
 - definition of, 3-20
 - how to use, 3-26
 - noncompliant products, 3-27
 - removing products from, 3-34
 - retrieving product information from, 3-33
 - tracking software dependencies, 3-27
 - volume label change, 3-33
- Product dependencies, 3-27
- Product description file
 - See PDF
- PRODUCT INSTALL command, 3-27
- Product list
 - obtaining, 3-10
 - VMSINSTAL.COM parameter, 3-10
- Product listing, 3-35
- Profiles
 - in MAIL, 6-38
 - in SYSMAN, 2-16
 - changing default directory, 2-17
 - changing privileges, 2-17
- Protected images, 16-11, 16-14, 16-15
- Protected subsystems, 6-16
 - enabling, 8-28
 - mounting volumes with, 8-15
- Protection
 - See also Protection codes
 - See also Security
 - ACL-based, 6-16, 6-32, 11-9
 - applying to public disk volumes, 8-15
 - applying to queues, 13-22 to 13-26
 - applying to system dump file, 15-4
 - assigning code when mounting a volume, 8-23
 - changing, 8-16
 - changing with PROTECTION qualifier, 8-18
 - default, 9-7
 - changing, 9-8
 - directory
 - changing with SET SECURITY /PROTECTION command, 9-12
 - specifying with CREATE/DIRECTORY command, 9-12
 - specifying with /PROTECTION qualifier, 9-12
 - disk file, 9-6
 - disk volumes, 8-16
 - displaying, 9-7
 - file, 9-4
 - default, 9-6, 9-7
 - directory, 9-4, 9-10
 - disk, 9-4
 - for public disk volumes, 8-15
 - for system dump file, 15-4
 - ISO 9660-formatted media, 8-17
 - magnetic tape, 9-12
 - for interchange environments, 8-20
 - format for object, 11-8
 - mask, 8-17
 - mounting a volume with protected subsystems, 8-28
 - of devices, 7-5
 - overriding, 8-25
 - qualifier for SET VOLUME command, 8-18
 - UIC-based codes, 6-16
 - volumes, 9-4
 - disk, 8-14, 8-15
 - ISO 9660-formatted media, 8-17
 - standard-labeled, 8-19
 - tape, 8-15, 8-19, 9-22
 - VOLPRO privilege, 8-14
 - volume set
 - ISO 9660-formatted media, 8-17

- Protection checks
 - MOUNT command
 - overriding, 8-23
- Protection codes
 - access types, 11-8
 - changing, 9-8
 - format, 11-8
 - null access specification, 11-8
 - specifying, 9-6
- Protocols
 - for using TCP/IP, 21-13
 - FTP, 21-14
 - LASTport, 23-5
 - LASTport/Disk, 23-5
 - LASTport/Tape, 23-6
 - RCP, 21-14
 - RLOGIN, 21-13
 - TELNET, 21-13
 - virtual terminal, 21-13
- Proxy accounts
 - adding, 6-35
- Proxy authorization files
 - NET\$PROXY.DAT, 6-4, 6-35
 - NETPROXY.DAT, 6-4, 6-35
- Proxy database
 - logical name defining location, 5-8
 - managed by security server, 5-5
- Proxy logins
 - controlling system use, 6-37
- PRTSMB symbiont, 13-3
 - on LAT printers, 13-78
- PTF (product text file), 3-20
- Public volumes
 - access to, 8-8
 - conditions for using, 8-8
 - creating with SYSTEM qualifier, 8-24
 - definition, 8-8
 - initializing, 8-12
 - guidelines, 8-13
 - load balancing, 8-9
 - mounting, 8-20, 8-21
 - in system startup, 5-11
 - mounting volume sets, 8-30
 - on large configurations, 8-8
 - on small configurations, 8-8
 - planning, 8-8
 - protecting, 8-15
 - setting protection, 8-18
 - testing disk performance, 8-8
- PURGE command
 - NCP, 21-11
 - saving disk space, 8-53

Q

- QMAN\$MASTER.DAT file
 - See Master file of queue database
- QMAN\$MASTER logical name, 12-6
 - defining during system startup, 5-8
 - defining to reduce system disk I/O, 16-8
- QUANTUM system parameter, 26-8
- Queue characteristics
 - canceling, 13-59
 - commands for managing, 13-57
 - defining, 13-58
 - definition, 13-28
 - deleting, 13-59
 - obtaining information about, 13-51, 13-59
 - problems
 - deleting, 13-82
 - mismatch, 13-81
 - sample use of, 13-28
 - specifying
 - on a job, 13-28
 - on a queue, 13-19, 13-28, 13-59
 - storage of in queue database, 13-17
- Queue commands
 - affected by multiple queue managers, 12-4
 - creating a queue database, 12-7
 - creating queues, 13-48
 - deleting queues, 13-57
 - displaying information about the queue
 - manager, 12-11
 - displaying jobs, 13-68
 - displaying queues, 13-49
 - enabling autostart, 13-16, 13-18, 13-48
 - managing banner pages, 13-60
 - managing characteristics, 13-57
 - managing device control libraries, 13-65
 - managing forms and stock, 13-61
 - managing queues, 13-46
 - modifying jobs, 13-70
 - modifying queues, 13-53
 - pausing queues, 13-53
 - relationship between starting and enabling
 - autostart, 13-4
 - setting UIC-based protection, 13-23
 - showing UIC-based protection, 13-23
 - specifying options, 13-19, 13-31
 - starting queues
 - autostart, 13-16, 13-18
 - nonautostart, 13-17, 13-48
 - starting the queue manager, 12-7
 - caution, 12-8
 - creating an additional queue manager, 12-10
 - restarting, 12-9
 - specifying failover list, 12-9
 - specifying name of queue manager, 12-10

Queue commands

- starting the queue manager (cont'd)
 - specifying nodes to run the queue manager, 12-6
- stopping
 - all queues on a node, 12-9
 - queues, 13-49, 13-54
 - the queue manager, 12-9

Queue configurations

- sample batch queuing system, 13-4 to 13-7
- sample output queuing environment, 13-8 to 13-14

Queue database

- See also Journal file of queue database
- See also Master file of queue database
- See also Queue file of queue database
- closing, 12-9
- creating, 12-7
- default location, 12-5
- definition, 12-4
- detection of queue corruption, 13-84
- determining location, 12-7
- determining location of master file, 12-5
- determining location of queue and journal file, 12-6
- files comprising, 12-4
- for multiple queue managers, 12-4, 12-5
 - naming convention, 12-5
- function, 12-4
- logical name defining location, 5-8
- managing, 12-1
- mounting the disk holding, 12-6
- moving, 12-5
- moving to reduce system disk I/O, 16-8
- requirement in a VMScluster environment, 12-5, 12-6
- restoring a damaged, 12-12
- saving, 12-11
- specifying location, 12-5

Queue failover, 13-4

Queue file of queue database, 12-4

- See also Queue database
- location, 12-6, 12-7
 - changing, 12-6
- saving, 12-12

Queue files

- detection of corruption, 13-84

Queue forms

- See Forms

Queue managers

- automatic restart, 12-3, 12-7, 12-9, 12-17
- availability of, 12-9, 12-17
- communication with job controller, 12-3
- creating an additional, 12-10
- default name, 12-4
- definition, 12-1
- displaying information about, 12-11

Queue managers (cont'd)

- failover, 12-3, 12-17
 - forcing, 12-9
 - list of nodes for, 12-9, 12-17
- function, 12-1
- improving performance, 12-13
- limiting nodes that can run, 12-9
- managing, 12-1
- moving to another node in a VMScluster system, 12-10
- multiple, 12-3, 12-10
 - commands affected by, 12-4
 - managing queues with, 12-10
 - moving a queue to another queue manager, 12-10
 - restriction, 12-3
 - specifying queue manager name, 12-4, 12-10
- naming, 12-10
- relationship to queues, 12-1
- restarting, 12-9
- role in queuing process, 12-2, 13-2
 - print jobs, 13-3
- role in starting active autostart queues, 13-4
- specifying name of, 12-10
- specifying preferred order in which nodes run, 12-9
- starting initially, 12-7
- stopping, 12-9
- troubleshooting, 12-14

Queue master file

- logical name defining location, 5-8

Queue names

- default for batch queue, 13-5
- default for print queue, 13-8
- defining, 13-15

Queue options, 13-18

- banner pages, 13-33
- characteristics, 13-28
- controlling job performance and resources, 13-28
- controlling page and line overflow, 13-43
- device control libraries, 13-44
- forms, 13-42
- qualifiers for specifying, 13-19 to 13-22
- restricting access, 13-22
- retaining jobs, 13-26

Queues

- See also Generic queues

- activating autostart, 13-15, 13-16, 13-49
- allowing jobs to complete before shutdown, 13-55
- assigning a default form, 13-63
- assigning device control libraries, 13-66
- autostart, 13-4
 - See also Autostart feature
 - See also Autostart queues
 - activating, 13-4

Queues

- autostart (cont'd)
 - on LAT queues, 13-4, 13-10
 - starting, 13-4
- availability of, 13-4, 13-15
- batch
 - setting up for vector processing, 26-7
- batch execution
 - See Batch queues
- canceling characteristics, 13-59
- changing DEFAULT form, 13-63
- changing options on, 13-53
- characteristics, 13-28
- closing, 13-53
- commands
 - See Queue commands
- creating, 13-15, 13-48
 - autostart execution, 13-15, 13-16
 - generic, 13-17
 - nonautostart execution, 13-16
- defining a characteristic, 13-58
- defining a form, 13-61
- deleting, 13-57
- deleting a job in, 13-74
- disabled, 13-84
- displaying information about, 13-49
- failover of, 13-15
- forms, 13-42
- gathering information with F\$GETQUI, 13-51
- generic batch, 13-7
 - See also Generic queues
- generic output, 13-11
 - See also Generic queues
- holding and releasing jobs, 13-71
- in a VMScluster environment, 20-2
- initializing, 13-48
- managing with multiple queue managers, 12-10
- merging, 13-56
- modifying, 13-53
- monitoring, 13-49
- mounting a form on, 13-42, 13-64
- moving jobs from one queue to another, 13-57
- moving to another queue manager, 12-10
- on standalone workstations
 - batch, 13-5
 - output, 13-8
- output execution
 - See Execution queues
 - See Output queues
- pausing, 13-53
- print
 - See Output queues
- problems deleting, 13-82
- protection, 13-22 to 13-26
- reinitializing existing, 13-53

Queues (cont'd)

- requeuing
 - executing job, 13-73
 - pending job, 13-73
- simplifying startup, 13-4
- spooling printers, 7-14, 13-14
- stalled
 - status of, 13-79, 13-81
 - troubleshooting, 13-81
- starting
 - autostart, 13-48
 - in startup command procedure, 13-18
 - in system startup, 5-12
 - nonautostart, 13-16, 13-48
- startup command procedure, 5-12
 - sample, 13-18
- stopped
 - status of, 13-79
- stopping, 13-54
 - abruptly, 13-54
 - all on a node, 13-55
 - before shutdown, 13-55
 - smoothly, 13-54
- types of, 13-2
 - output execution, 13-3
- Queue status, 13-49, 13-50
 - determining, 13-78
- Queuing system
 - See Batch and print queuing system
- QUOTA.SYS file
 - See Quota file
- Quota file
 - Analyze/Disk_Structure utility creates version, A-9
 - contents, 8-49
 - creating, 8-51
 - deleting, 8-53
 - disabling, 8-53
 - QUOTA.SYS, A-9
 - requirements for, 8-51
 - reserved file, A-9
 - UIC [0,0] in, 8-49
 - updating, 8-53
- Quotas
 - See also Process limits
 - See also UAFs (user authorization files), resource limits
 - disk
 - See Disk quotas
 - required to run UETP, 17-23

R

- RAID (redundant arrays of independent disks)
 - volume shadowing support, 10-40

- RCP, 21-14
- Read access
 - See also Access types, read
 - for disk directory files, 9-11
 - for disk files, 9-6
 - granting through protection codes, 11-8
- Read error
 - if returned when booting, 4-16
- Read operation
 - See Access types, read
- Read/write disk
 - partitions, 23-2
- Real-time priority, 6-31
- Rebuilding volumes, 7-4
- Reconfiguration of software installation options, 3-32
- Records
 - blocking on tapes, 8-6
 - size, 8-26
- Recovering lost files, 8-57
- Reducing I/O on system disk, 16-8
- Redundant arrays of independent disks
 - See RAID
- Registering images with system version
 - dependencies, 5-22
- Reinitializing volumes, 8-45
- Release notes, extracting from software products, 3-29
- release notes on software products, 3-29
- Release Notes option
 - VMSINSTAL.COM, 3-17
- Releasing a job, 13-71
- Remote directory listings, 21-14
- Remote file access, 21-14
- Remote identifiers, 11-11
- Remote InfoServer device
 - BIND command, 23-14
 - making available, 23-14
 - mounting during system startup, 23-14
- Remote login, 21-13
- Remote monitoring
 - limitation, 18-43
 - mixed-version VMScluster systems, 18-43
- Remote node database
 - logical name defining location, 5-8
- Remote nodes
 - definition, 2-13
- Remote printing, 13-13
- Remote queue status, 13-51
- Remote terminals, 7-11
- Remote terminal service, 21-13
- Removing installed software products, 3-34
- Repairing disk structure errors, 8-56
- REPLY command
 - canceling a user request, 8-27, 18-19
 - closing current operator log file, 18-17
 - disabling operator terminals, 2-24, 18-18
- REPLY command (cont'd)
 - enabling operator terminals, 2-23, 18-17
 - specifying terminal, 18-18
 - ensuring that correct volume is mounted, 8-41
 - initializing tape, 8-41
 - linking continuation volume to volume set, 8-41
 - opening a new operator log file, 18-17, 18-22
 - putting request in wait state, 8-27
 - replying to requests, 2-25, 8-27
 - response not recorded in operator log file, 18-19
 - sending messages to users, 2-22
- REPLY/ENABLE=SECURITY command
 - enabling security operator terminals, 18-29
- Reporting errors
 - disk structure, 8-56
- Reporting on product dependencies, 3-27
- Reports
 - AUTOGEN, 14-22
 - SHOW CLUSTER
 - adding classes of data, 20-4
 - adding data, 20-8
 - adding fields of data, 20-4
 - changing default at startup, 20-14
 - command to modify, 20-10
 - compressing reports, 20-11
 - controlling displays, 20-5, 20-10
 - controlling with command procedures, 20-5, 20-16
 - moving display, 20-12
 - organization of, 20-4
 - panning, 20-6
 - scrolling, 20-14
- REQUEST command
 - recording request in operator log file, 18-19
 - sending requests to an operator, 2-25, 18-19
- Request identification number
 - in operator log file, 18-19
- Requeuing
 - executing job, 13-73
 - pending job, 13-73
- Reserved files, A-4
 - backup log file (BACKUP.SYS), A-9
 - bad block file (BADBLK.SYS), A-8
 - continuation file (CONTIN.SYS), A-9
 - index file (INDEXF.SYS), A-5
 - list of, 8-4
 - master file directory (MFD), A-8
 - pending bad block log file (BADLOG.SYS), A-9
 - quota file (QUOTA.SYS), A-9
 - storage bit map file (BITMAP.SYS), A-8
 - volume security profile (SECURITY.SYS), A-9
 - volume set list file (VOLSET.SYS), A-9
- Reset button on POLYCENTER Software
 - Installation utility, 3-40

- Reset modules, 13-44
 - See also Device control modules
 - Resetting saved information, 3-40
 - Resident images
 - installing (Alpha), 16-11, 16-12, 16-13
 - in system startup, 5-13
 - Resource accounting, 19-1
 - Resource limits
 - See Process limits
 - Resource use
 - producing reports of, 19-4
 - Restarting OPCOM (operator communication manager), 2-22
 - Restarting the queue manager, 12-9
 - Restore operations
 - with BACKUP, 10-28
 - Restoring directories, 10-29
 - Restoring files, 10-28
 - from an image backup, 10-42
 - from an incremental backup, 10-44
 - levels of directory access, 10-30
 - Restoring the queue database, 12-12
 - Restricted accounts, 6-29
 - RESTUSER.COM command procedure, 10-36
 - Resuming printing of an output job, 13-75, 13-76, 13-77
 - Resuming queue status, 13-51
 - Retained job status, 13-74
 - definition, 13-69
 - releasing jobs in, 13-71
 - Retaining jobs in a queue
 - changing retention of a job, 13-74
 - Rights databases, 5-8
 - RIGHTSLIST.DAT files
 - default protection, 6-5
 - moving to reduce system disk I/O, 16-8
 - RIGHTSLIST logical name
 - defining during system startup, 5-8
 - defining to reduce system disk I/O, 16-8
 - RLOGIN, 21-13
 - RMS
 - access to files at record level, 9-13
 - RMS_EXTEND_SIZE system parameter, 16-7
 - Root directory
 - adding to an existing system disk, 2-31
 - Root volume
 - definition, 8-31
 - MFD on, 8-31
 - Routers
 - Level 1
 - definition, 21-5
 - Level 2
 - definition, 21-6
 - network
 - definition, 21-5
 - end node, 21-6
 - Routing
 - levels of, 21-5
 - network
 - adaptive, 21-5
 - area, 21-6
 - definition, 21-5
 - Routing nodes
 - See Routers
 - RSM (Remote System Manager)
 - startup restrictions, 23-12
 - RT-11 volume
 - block-addressable, 9-24
 - RTAN devices, 7-11
 - RUN command
 - effect of installed images on, 16-10
 - RV60 optical disk drives
 - supported by UETP, 17-7
 - RVN (relative volume number)
 - See File identification, relative volume number
- ## S
-
- Satellites
 - booting, 22-20
 - cross-architecture booting, 22-23
 - LAN MOP and DECnet MOP downline load service, 22-20
 - migrating to LAN MOP downline load service, 22-20
 - SATELLITE_PAGE.COM command procedure, 5-6, 15-20
 - execution of during system startup, 5-5
 - SAVEDUMP system parameter, 15-2, 15-3, 15-7
 - Save sets, 10-6, 10-23
 - Files-11 disk, 10-7
 - Get Save Set
 - VMSINSTAL.COM option, 3-15
 - listing the contents of, 10-18
 - magnetic tape, 10-6
 - multivolume, 10-29
 - name, 10-24
 - network, 10-7
 - on multiple tapes, 10-21
 - product
 - storing temporarily during installation, 3-15
 - sequential disk, 10-7
 - types, 10-6
 - Saving the queue database, 12-11
 - Scalars
 - definition, 26-4
 - SCBs (storage control blocks)
 - in storage bit map file, A-8
 - Scheduling, 6-2, 6-31
 - See also Priority, job scheduling
 - of batch jobs, 13-72
 - of print jobs, 13-32, 13-72

- SCROLL command, 20-13
- SCROLL keypad function, 20-7
- SCSI (Small Computer Systems Interface), 10-40, 21-9
- SCSNODE system parameter, 23-10
 - in DECdtm Services, 25-22
 - unique value for DECdtm use, 25-24
- SDA (System Dump Analyzer utility)
 - See System Dump Analyzer utility
- SDA CLUE commands
 - analyzing the system dump file, 15-2
 - saving contents of system dump file, 15-15
 - using to archive dump file information, 15-11
 - using to collect dump file information, 15-11
- Search lists
 - priority of installed images, 16-10
- Search path of Help Message database files, 5-28
- Secondary Bootstrap image
 - role in boot process, 4-2
- SECONDARY day
 - defining for accounts, 6-27
- Secondary page and swap files, 15-5
 - See also Page files
 - See also Swap files
 - creating
 - in system startup, 15-20
 - with AUTOGEN, 15-17, 15-24
 - with SYPAGSWPFILES.COM procedure, 15-20
 - improving system performance, 16-8
 - installing
 - in system startup, 15-20
 - requirement, 15-5
 - with SYPAGSWPFILES.COM command procedure, 15-20
 - installing during system startup, 5-6
- Secondary processors, 26-2
- Sections
 - global, 16-12
 - pages, 16-12
- Sectors
 - Files-11
 - definition, A-2
- Security
 - See also Protection
 - alarm messages, 18-21
 - alarms
 - enabling, 18-29
 - auditing
 - description, 18-25
 - enabling operator terminal, 18-29
 - audit log files
 - See also Security audit log files
 - authenticating user's identity, 21-14
 - enabling operator terminals, 18-21
 - messages in operator log file, 18-21
 - network
 - Security
 - network (cont'd)
 - See also DECnet, security providing for, 21-12
 - OPCOM alarm messages, 18-21
 - password management, 11-2
 - protected subsystems, 8-28
 - protecting public disk volumes, 8-15
 - protecting queues, 13-22
 - protecting the system dump file, 15-4
 - risk of compromise by installing images with privileges, 16-13
 - specifying alarm events, 18-21
 - VMScluster, 20-16
- SECURITY.AUDIT\$JOURNAL file
 - See Security audit log files
- SECURITY.SYS file
 - See Volume security profile
- Security auditing, 11-13
 - defining log file in system startup, 5-9
 - description, 18-25
 - disabling events, 18-28
 - displaying using SHOW AUDIT command, 18-27
 - enabling operator terminal, 18-29
- Security audit log files
 - closing, 18-30
 - creating new version, 18-30
 - definition, 18-2
 - reviewing, 18-26
- Security management, 6-16
 - in SYSMAN on remote nodes, 2-16
 - protecting queues, 13-22 to 13-26
- Security operator terminals, 18-29
- Security server
 - starting, 5-5
- Security Server process
 - intrusion database, 11-7
 - network proxy database, 11-7
- SECURITY_AUDIT.AUDIT\$JOURNAL file
 - See also Security audit log files
 - moving to reduce system disk I/O, 16-8
- SECURITY_SERVER
 - starting, 5-5
- SELECT command
 - in SHOW CLUSTER, 20-13
- Selective dump, 15-3
 - compared to physical dump, 15-3, 15-10
 - storing, 15-10
- Sending comments to Digital writers, iii
- SEQ (file sequence number)
 - See File identification, file sequence number
- Sequential-disk save set, 10-7
 - initializing, 10-8
 - mounting, 10-8

- Server queues, 13-3
- Server queue status, 13-51
- Services
 - bindings, 23-13
 - nodes, 24-4
 - password protection, 23-13
 - write protection, 23-13
- Sessions
 - maintaining on more than one terminal, 7-11
 - maintaining when disconnecting a terminal, 7-11
- SET (Field) command
 - SHOW CLUSTER, 20-11
- SET ACCOUNTING command
 - controlling which resources are tracked, 19-3
 - starting up a new accounting file, 19-3
- SET ACL command
 - for vector capability (VAX), 26-8
 - modifying disk file characteristics, 9-8
 - modifying file characteristics, 9-9
- SET AUDIT command
 - creating new version of security audit log file, 18-30
 - on local node only, 18-30
 - enabling security alarms, 18-29
 - to create new version of security audit log file, 18-30
 - to enable security auditing, 18-25
- SET command
 - in conversational boot, 4-6
 - NCP, 21-11
- SET DEVICE command
 - spooling printers, 7-14
- SET DIRECTORY command
 - changing directory characteristics, 9-12
 - limiting disk space consumed, 8-8
 - modifying disk file characteristics, 9-8
- SET ENTRY command, 13-70
 - changing forms, 13-80
 - changing scheduling priority, 13-72
 - holding jobs, 13-71
 - releasing jobs, 13-71
 - requeuing pending jobs, 13-73
 - specifying job retention, 13-74
- SET ENVIRONMENT command, 2-14
- SET FILE command
 - example, 6-33
 - modifying disk file characteristics, 9-8
 - using to assign an alias, 9-10
 - using to modify file characteristics, 9-10
- SET FUNCTION command, 20-7
 - in SHOW CLUSTER, 20-13
- SET HOST command, 21-13
- SET HOST/LAT command, 24-2
- SET LOGINS command, 16-4
- SET MAGTAPE command, 7-16, 8-42
- SETPARAMS.DAT file, 14-18
- SET PRINTER command, 7-13, 13-14
 - in system startup, 5-12
- SET PROFILE command
 - in SYSMAN, 2-17
- SET QUEUE command, 13-53
 - assigning a default form, 13-63
 - assigning characteristics, 13-59, 13-81
 - canceling characteristics, 13-59
 - controlling page overflow, 13-44
 - mounting a form, 13-64
 - setting block limits, 13-32
 - setting scheduling policy, 13-32
 - setting UIC-based protection on queues, 13-23
 - specifying banner pages, 13-60
 - specifying job processing options, 13-31
 - specifying queue options with, 13-19
 - specifying reset modules, 13-65
- SET QUORUM command
 - avoiding use of /CLUSTER with SYSMAN DO command, 20-21
- SET SECURITY command
 - for queues, 13-23
 - modifying file characteristics, 9-9
- SET SECURITY/PROTECTION command, 9-9
 - changing directory protection, 9-12
 - setting default protection, 9-8
- SET/STARTUP command
 - in conversational boot, 4-12
- SET TERMINAL command, 7-9, 13-14
 - despooling a printer, 7-15
 - determining characteristics of a LAT line, 7-12
 - enabling virtual terminals, 7-11
 - in system startup, 5-12, 7-10, 7-13
 - setting printer characteristics, 7-13
- SET TIMEOUT command, 2-18
- Setting system parameters
 - See Changing system parameters
- Setting up printers, 7-9, 7-13, 13-14
 - in system startup, 5-12, 7-13
 - LAT, 13-14
- Setting up terminals, 7-9, 7-13, 13-14
 - in system startup, 5-12, 7-10
 - using system parameters, 7-10
- Setup module, 13-44
 - See also Device control modules
 - specifying in forms, 13-42
- SET VOLUME command
 - changing protection code, 8-18
 - disabling high-water marking, 16-7
 - encoding labels on volumes, 8-29
 - modifying disk volume characteristics, 8-29
 - modifying file characteristics, 9-8
 - performing data checks, 8-29
 - specifying file retention periods, 8-54

- Shadow sets
 - backing up, 10-40
 - mounting disks, 10-42
 - restoring, 10-47
- Shareable images, 16-11
 - assigning logical names for, 16-17
- Shareable linkage sections, 16-13
- Shared resource
 - definition, 20-1
- Sheet-feed paper
 - specifying in forms, 13-42
- Short report format
 - See Console report during UETP
- SHOW ACCOUNTING command, 19-2
- SHOW ACL command, 26-8
- SHOW AUDIT command
 - to display event classes being audited, 18-27
- SHOW CLUSTER command
 - See Show Cluster utility
- Show Cluster utility (SHOW CLUSTER)
 - controlling displays, 20-5
 - defining keypad keys, 20-7
 - refreshing the screen, 20-10
 - reports, 20-4
 - controlling displays, 20-10
 - formatting, 20-5, 20-11, 20-14
 - initialization file, 20-10
 - using, 20-4
- SHOW command, 9-1
 - in conversational boot, 4-6
 - in SYSGEN, 14-31
 - specifying languages, 5-43, 5-44
 - time and date formats, 5-42
- SHOW CPU command, 26-3, 26-9
- SHOW DEVICES command, 7-2, 7-16, 9-5
 - checking mounted volumes, 8-37
 - determining the status of files, 8-44, 16-18
 - devices not shown, 7-7
 - examples, 7-2
 - for ISO-9660 formatted devices, 7-5
 - volume rebuild status, 7-4
- SHOW ENTRY command, 13-68
- SHOW ENVIRONMENT command, 2-14
- Showing system parameters
 - with conversational boot, 4-6
 - with SYSGEN, 14-31
 - with SYSMAN, 14-27
- SHOW INTRUSION command, 11-7
- SHOW MEMORY command
 - determining size of page and swap files, 15-23
 - displaying page and swap files, 15-6
 - monitoring page file usage, 15-6
 - monitoring swap file usage, 15-10
 - showing the size of nonpaged pool, 26-7
- SHOW PROCESS command, 9-5, 26-9
- SHOW PROFILE command
 - in SYSMAN, 2-17
- SHOW PROTECTION command, 9-5
- SHOW QUEUE command
 - showing all jobs, 13-50
 - showing batch jobs, 13-50
 - showing brief information, 13-50
 - showing complete information, 13-50
 - showing completion status for jobs, 13-26
 - showing files associated with a job, 13-50
 - showing generic queues, 13-50
 - showing jobs of a specified status, 13-50
 - showing output execution queues, 13-50
 - showing queue status, 13-49
 - showing total number of jobs, 13-50
- SHOW QUEUE/MANAGER command, 12-11
- SHOW QUOTA command, 8-52
- SHOW SECURITY command
 - for queues, 13-23
- Show software products, 3-35
- SHOW/STARTUP command
 - in conversational boot, 4-12
- SHOW_CLUSTER\$INIT.COM command
 - procedure, 20-14
- SHOW_CLUSTER\$INIT logical name, 20-14
- Shutdown
 - See SHUTDOWN.COM command procedure
 - See System shutdown
- SHUTDOWN\$DECNET_MINUTES logical name, 4-33
- SHUTDOWN\$DISABLE_AUTOSTART logical name, 4-33, 13-56
- SHUTDOWN\$INFORM_NODES logical name, 4-33
- SHUTDOWN\$MINIMUM_MINUTES logical name, 4-33
- SHUTDOWN\$QUEUE_MINUTES logical name, 4-33
- SHUTDOWN\$TIME logical name, 4-33
- SHUTDOWN\$VERBOSE logical name, 4-33
- SHUTDOWN.COM command procedure, 4-27
 - See also System shutdown
 - customizing, 4-32, 4-33
 - defining time before shutdown, 4-33
 - example, 4-30
 - executing with SYSMAN, 4-34
 - how to use, 4-28
 - options
 - time of shutdown, 4-33
 - order of events, 4-31
 - reboot options, 4-29
 - required privileges, 4-27
 - when to use, 4-27
- Site-independent startup command procedure
 - See STARTUP.COM command procedure

- Site-specific products
 - startup database, 5-19
- Site-specific shutdown command procedure
 - SYSHUTDWN.COM, 4-27
- Site-specific startup command procedure
 - order of execution, 4-4
 - SATELLITE_PAGE.COM
 - See SATELLITE_PAGE.COM command procedure
 - SYCONFIG.COM
 - See SYCONFIG.COM command procedure
 - SYLOGICALS.COM
 - See SYLOGICALS.COM command procedure
 - SYPAAGSWPFILES.COM
 - See SYPAAGSWPFILES.COM command procedure
 - SYSECURITY.COM
 - See SYSECURITY.COM command procedure
 - SYSTARTUP_VMS.COM
 - See SYSTARTUP_VMS.COM command procedure
- Slicing images, 16-12
- Small computer systems interface
 - See SCSI
- SMISERVER process, 2-12
 - attributes of, 2-16
 - changing, 2-16
 - starting, 2-12
 - in system startup, 5-5
- SMP (symmetric multiprocessing)
 - See Multiprocessing
- SMP_CPUS system parameter, 26-3, 26-6
- Snapshot facility
 - concepts, 4-19
 - preparing system for, 4-20
 - preparing system startup files for, 4-20
 - required process limits, 4-20
 - supported applications, 4-20
 - when using with DECwindows, 4-21
- Software errors
 - OPCOM failure, 2-22
 - queue manager, 12-14
 - when booting, 4-16
- Software installation
 - See Installing software
 - identifying kit location, 3-28
- Software kits
 - See Software products
- Software license
 - definition, 3-8
- Software Performance Reports
 - See SPRs

- Software products
 - See also Installing software
 - consistent installation of, 3-22
 - converting formats, 3-33
 - copying to new locations, 3-33
 - dependencies on other products, 3-27
 - installing more than one, 3-29
 - reconfiguring after installing, 3-32
 - registering noncompliant software, 3-27
 - release notes, 3-29
 - removing after installing, 3-34
 - tracking dependencies of, 3-27
- SOGW user category abbreviation, 11-8
- Sort/Merge utility (SORT/MERGE)
 - optimizing batch queues for, 13-31
- Source and destination locations of software kits, 3-28
- Source files for installing software
 - PCSI\$SOURCE location, 3-28
 - specifying location, 3-28
- Source parameter
 - VMSINSTAL.COM, 3-11
- SPAWN function
 - LAN Control Program (LANCP) utility, 22-9
- Spawn mode
 - as execution mode for a startup procedure, 5-19
- Spooled printers
 - See also Printers, spooled
 - See Printers, spooled
- SPRs (Software Performance Reports)
 - including system dump file with, 15-2, 15-11
 - submitting to report system failure, 15-11
- STABACKIT.COM command procedure, 10-52, 10-53
- Stalled job status, 13-69
- Stalled queues
 - status, 13-51
- Standalone BACKUP
 - booting, 10-52, 10-54
 - building, 10-51, 10-53
 - definition, 10-51
 - qualifiers, 10-50
 - relation to Backup utility, 10-51
 - using to back up the system disk, 10-51, 10-55, 10-59
 - using to restore the system disk, 10-57
- START/CPU command, 26-3, 26-6
- Starting InfoServer Client for OpenVMS, 23-10
- Starting queues, 13-4
 - autostart, 13-48
 - in startup command procedure, 13-18
 - relationship to activating an autostart queue, 13-4
 - nonautostart, 13-16, 13-48
 - in startup command procedure, 13-18

- Starting queue status, 13-51
- Starting the LAT software
 - with LAT\$STARTUP.COM, 5-15, 24-10, 24-11
- Starting the queue manager, 12-7
 - initially, 12-7
 - restarting, 12-9
- STARTNET.COM command procedure, 5-15, 7-8
- START/QUEUE command, 13-18
 - activating an autostart queue, 13-49
 - assigning a default form, 13-63
 - assigning characteristics, 13-59
 - canceling characteristics, 13-59
 - controlling page overflow, 13-44
 - mounting a form, 13-64
 - resuming printing of a suspended job, 13-76
 - setting block limits, 13-32
 - setting scheduling policy, 13-32
 - setting UIC-based protection on queues, 13-23
 - specifying autostart information, 13-15
 - specifying banner pages, 13-60
 - specifying job processing options, 13-31
 - specifying queue options with, 13-19
 - specifying reset modules, 13-65
 - starting a generic queue, 13-17
 - starting a nonautostart queue, 13-48
- START/QUEUE/MANAGER command, 12-7, 12-9
 - caution, 12-8
 - creating an additional queue manager, 12-10
 - creating a queue database, 12-7
 - specifying failover list, 12-9
 - specifying name of queue manager, 12-10
 - specifying nodes to run the queue manager, 12-6
 - storage of, 12-7
- STARTUP\$AUTOCONFIGURE_ALL symbol, 7-8
- STARTUP\$INTERACTIVE_LOGINS symbol, 5-16
- STARTUP\$STARTUP_LAYERED logical name, 5-18
- STARTUP\$STARTUP_VMS logical name, 5-18
- STARTUP.COM command procedure, 4-4, 5-3
 - configuring devices, 5-7, 7-6
 - definition, 5-2
 - description, 4-12
 - if it does not execute, 4-16
 - message indicating execution of, 4-5
 - tasks performed by, 4-4, 4-12, 16-10
- STARTUP command, 2-11
 - See also Startup database in SYSMAN, 5-18
- Startup command procedure
 - See also Site-specific startup command procedure
 - See also STARTUP.COM command procedure
 - booting without, 4-8
 - changing execution mode, 5-21

- Startup command procedure (cont'd)
 - changing node restrictions, 5-21
 - changing startup phase, 5-21
 - creating your own, 7-10, 7-13
 - enabling a temporarily disabled, 5-22
 - known file lists, 5-12
 - modifiable
 - See Site-specific startup command procedure
 - node restriction, 5-19
 - passing parameters to, 5-19
 - preventing from executing, 5-21
 - temporarily, 5-22
 - required
 - See STARTUP.COM command procedure
 - SATELLITE_PAGE.COM
 - See SATELLITE_PAGE.COM command procedure
 - setting up output devices, 13-14
 - site-independent
 - See also STARTUP.COM command procedure
 - specifying an alternate, 4-12
 - as the default, 4-13
 - site-specific, 5-2, 5-3
 - See also Site-specific startup command procedure
 - announcements, 5-14
 - .COM version, 5-3
 - creating your own, 5-2
 - definition, 5-2
 - modifying, 5-4
 - modifying to perform site-specific operations, 5-2
 - order of execution, 5-3
 - required location, 5-4
 - .TEMPLATE version, 5-3
 - use in VMSKITBLD, 5-3
 - versions of, 5-3
 - specifying execution mode, 5-20
 - specifying node restrictions, 5-20
 - specifying startup phase, 5-20
 - starting queues, 13-18
- SYCONFIG.COM
 - See SYCONFIG.COM command procedure
- SYLOGICALS.COM
 - See SYLOGICALS.COM command procedure
- SYPAWSWPFILES.COM
 - See SYPAWSWPFILES.COM command procedure
- SYSECURITY.COM
 - See SYSECURITY.COM command procedure
- when errors prevent you from logging in, 4-8

- Startup database
 - adding files to, 5-20
 - changing information in, 5-21
 - definition, 5-18
 - deleting records in, 5-21
 - disabling files in, 5-22
 - reenabling disabled files in, 5-22
 - restriction, 5-21
 - showing contents of, 5-20
 - showing name of the target, 5-20
 - specifying the current, 5-20
- Startup phases
 - determining order of, 5-18
 - layered product, 5-18
 - END, 5-19
 - LPBEGIN, 5-19
 - LPBETA, 5-19
 - LPMAIN, 5-19
 - specifying, 5-20
 - operating system
 - BASEENVIRON, 5-4
 - CONFIGURE, 5-4
 - DEVICE, 5-4
 - INITIAL, 5-4
- Startup restrictions
 - InfoServer Client for OpenVMS software, 23-12
 - PATHWORKS, 23-12
 - RSM, 23-12
 - SYSMAN, 23-12
- STARTUP SET OPTIONS command, 4-15
- STARTUP SHOW OPTIONS command, 4-16
- STARTUP_P1 system parameter, 4-14
- STARTUP_P2 system parameter, 4-14
 - SYSMAN startup logging, 4-15
- Status of jobs
 - See Job status
- Status of queues
 - See Queue status
- Status of volume rebuilds, 7-4
- \$STATUS values
 - accessing for uninstalled messages, 5-26
- Stock
 - See also Forms
 - commands used with, 13-61
 - mismatch, 13-42
 - troubleshooting, 13-80
 - specifying, 13-42
- STOP/CPU command, 26-3, 26-6
- Stopped queue status, 13-51
- Stop pending queue status, 13-51
- Stopping queue
 - status, 13-51
- Stopping queues, 13-54
 - abruptly, 13-54
 - all queues on a node, 12-9, 13-55
 - smoothly, 13-54
- Stopping the queue manager, 12-9
- STOP/QUEUE command, 13-53, 13-75
- STOP/QUEUE/MANAGER/CLUSTER command, 12-9
- STOP/QUEUE/NEXT command, 13-49, 13-54
 - with autostart queues, 13-54
- STOP/QUEUE/RESET command, 13-49, 13-54
 - with autostart queues, 13-54
- STOP/QUEUES/ON_NODE command, 12-9
 - entering before shutting down a system, 13-55
 - relationship to DISABLE AUTOSTART /QUEUES command, 13-55
- Storage bit map file
 - BITMAP.SYS, A-8
 - definition, A-8
 - reserved file, A-8
 - storage control block (SCB), A-8
- Storage control blocks
 - See SCBs
- StorageWorks RAID Array, 10-42
- SUBMIT command
 - preventing users from executing, 13-53
 - processing of, 13-2
 - specifying characteristics, 13-28
 - specifying job retention, 13-74
 - specifying scheduling priority, 13-72
- SUBMON.COM command procedure
 - sample, 18-41
 - used with Monitor utility, 18-39
- Subprocesses
 - creating with the LANCP SPAWN function, 22-9
 - subprocess creation limit, 6-2, 6-43
- Substituting volumes, 8-27
- Subsystem ACEs
 - example, 11-10
- Subsystems
 - protected, 8-28
- Supervisor mode
 - logical names, 16-15
- Suppressing message display, 14-12
- Suspending a job, 13-75
- SWAPFILE.SYS file, 15-4
 - See also Swap files
- SWAPFILE_n_NAME symbol, 15-17, 15-24
- SWAPFILE_n_SIZE symbol, 15-18, 15-24
- Swap files
 - changing sizes
 - with SWAPFILES.COM, 15-25
 - creating
 - with AUTOGEN, 15-17, 15-24
 - with SYSGEN, 15-18
 - definition, 15-4
 - deleting after creating a new version, 15-28
 - displaying, 15-6
 - displaying the size calculated by AUTOGEN, 15-17, 15-22

- Swap files (cont'd)
 - fragmentation of, 15-27
 - installing
 - in system startup, 5-5, 5-6, 15-5, 15-19, 15-21
 - when resized with AUTOGEN, 15-18, 15-22
 - with SYPAGSWPFILES.COM procedure, 15-20
 - with SYSGEN, 15-19
 - location
 - specifying for individual files, 15-24
 - message indicating high fragmentation, 15-27
 - monitoring usage of, 15-6, 15-10
 - mounting disk during system startup, 5-6, 15-20
 - moving to improve performance, 16-8
 - on a satellite, 15-20
 - primary, 15-5
 - purging, 15-28
 - requirements
 - location, 15-5
 - secondary, 15-5, 15-20
 - size
 - See Swap file sizes
 - tasks for managing, 15-1
 - VMScluster satellite node, 5-6
- SWAPFILES.COM command procedure
 - changing size of primary page, swap, and dump files, 15-25
- Swap file sizes
 - calculating
 - manually, 15-9
 - with AUTOGEN, 15-5
 - changing
 - recommended method, 15-22
 - when to increase, 15-10
 - with AUTOGEN, 15-22
 - with SYSGEN, 15-26
 - determining current, 15-23
 - displaying AUTOGEN's calculations, 15-17, 15-22
 - specifying
 - for individual files, 15-22, 15-24
 - total for multiple files, 15-22, 15-23
 - when to increase, 15-10
- SWAPFILE symbol, 15-24
- Swapping
 - to move information between physical memory and files stored on disk, 15-4
- SYCONFIG.COM command procedure, 5-3
 - AUTOGEN failure, 7-9
 - configuring devices, 7-6
 - in startup, 5-4
 - modifying to connect special devices, 5-7
 - modifying to mount disks early, 5-12
 - STARTUP\$AUTOCONFIGURE_ALL symbol, 7-8
- SYLOGICALS.COM command procedure, 5-3, 5-8
 - in system startup, 5-5
 - mounting the queue database disk, 12-6
 - redefining location of master file of queue database, 12-6
 - redefining location of system files, 16-8
 - to specify default state of operator log files, 18-23
- SYLOGIN.COM command procedure, 5-16, 6-20
 - ensuring execution, 6-13
 - sample systemwide, 6-20
- SYLOGOUT.COM command procedure, 6-22
- Symbionts, 13-53
 - bypass formatting, 13-44
 - communicating with, 13-75
 - default, 13-3
 - determining, 13-78
 - for LAT printers, 13-3, 13-78
 - function of, 12-3, 13-2
 - LATSYM, 13-3, 13-78
 - PRTSMB on LAT printers, 13-78
 - role in processing print jobs, 13-3
 - user-written, 13-3
- Symbols
 - See also Logical names
 - defining in MODPARAMS.DAT file, 14-19, 15-17
 - for page, swap, and dump file sizes, 15-23 to 15-25
 - for system parameters, 14-19
 - NUM_ETHERADAPT, 14-21
 - NUM_NODES, 14-21
 - PAGEFILE_n_NAME, 15-17
 - PAGEFILE_n_SIZE, 15-18
 - STARTUP\$AUTOCONFIGURE_ALL, 7-8
 - STARTUP\$INTERACTIVE_LOGINS, 5-16
 - SWAPFILE_n_NAME, 15-17
 - SWAPFILE_n_SIZE, 15-18
- Symmetric multiprocessing
 - See Multiprocessing
- Symmetric vector processing configuration, 26-4
- SYPAGSWPFILES.COM command procedure, 5-3, 15-19, 15-20
 - execution of during system startup, 5-5, 5-6
 - modifying to install page and swap files, 5-6
- SYS\$ANNOUNCE logical name, 5-14
- SYS\$AUDIT_SERVER_INHIBIT logical name, 5-9, 18-27
- SYS\$BATCH default queue name, 13-5
- SYS\$DECDTM_INHIBIT logical name, 25-21
- SYS\$ERRORLOG logical name
 - defining during system startup, 5-8
 - defining to reduce system disk I/O, 16-8
- SYS\$INPUT logical name, 17-31

SYS\$JOURNAL logical name, 25-3
 SYS\$LANGUAGE logical, 5-48
 SYS\$MONITOR logical name
 defining during system startup, 5-8
 defining to reduce system disk I/O, 16-8
 SYS\$OUTPUT logical name, 17-33
 SYS\$PRINT default queue name, 13-8
 SYS\$QUEUE_MANAGER.QMAN\$JOURNAL file
 See Journal file of queue database
 SYS\$QUEUE_MANAGER.QMAN\$QUEUES file
 See Queue file of queue database
 SYS\$QUEUE_MANAGER queue manager
 as default queue manager, 12-4
 SYS\$STARTUP logical name, 4-4, 5-2, 5-18
 SYS\$SYLOGIN logical name, 5-17
 SYS\$SYSTEM:NETNODE_REMOTE.DAT file
 changing location of, 21-10
 contains configuration database, 21-10
 SYS\$TEST logical name, 17-4, 17-11, 17-20
 SYS\$TIMEZONE_DIFFERENTIAL logical name,
 5-35
 SYS\$UPDATE logical name
 with VMSINSTAL.COM command procedure,
 3-9
 SYS\$WELCOME logical name, 5-15
 SYSBOOT
 See Conversational boot
 SYSBOOT.EXE file, 4-2
 SYSDUMP.DMP file, 15-2
 See also System dump files
 protection, 15-4
 required location, 15-3
 SYSECURITY.COM command procedure, 5-3,
 5-9
 during system startup, 5-5
 SYSGEN (System Generation utility)
 See System Generation utility
 SYSGEN parameters
 See System parameters
 SYS\$SHUTDWN.COM command procedure, 4-27
 SYS\$LOST.DIR file
 lost files in, 8-57
 SYSMAN
 See System Management utility
 SYSMANINI logical name, 2-19
 SYSMAN parameters
 See System parameters
 SYS\$PRV privilege
 giving rights of system user, 11-8
 SYSTARTUP_VMS.COM command procedure,
 5-3 to 5-17
 creating systemwide announcements, 5-14
 defining announcements in, 5-14
 defining location of systemwide login procedure,
 5-17
 defining welcome messages in, 5-15

SYSTARTUP_VMS.COM command procedure
 (cont'd)
 disabling error checking in, 5-10
 editing to start DECnet for OpenVMS VAX,
 5-15
 enabling autostart, 5-12
 freeing page file of dump information, 15-4,
 15-17
 installing known images, 5-12, 16-10
 installing resident images (Alpha), 5-13
 invoking LAT command procedures, 5-15
 invoking the System Dump Analyzer utility,
 15-15
 limiting the number of interactive users, 5-16
 making remote InfoServer disks available,
 5-13, 23-14
 message indicating execution of, 4-5
 modifying to perform site-specific operations
 during system startup, 5-10
 operations performed in, 5-10
 purging the operator log file, 5-14
 saving contents of system dump file in, 15-15
 setting printer device characteristics, 5-12,
 7-13
 setting terminal device characteristics, 5-12,
 7-10
 special consideration about operator assistance
 for MOUNT command, 5-11
 starting InfoServer Client for OpenVMS, 5-13,
 23-10
 starting queues, 5-12
 submitting batch jobs from, 5-14
 SYSTEM account
 changing passwords for security of, 2-9
 exercising caution with privileges, 2-9
 initial modification, 6-9
 in UAF, 6-6, 6-7
 logging in to, 2-10
 setting process quotas for efficient backups,
 10-9
 using AUTHORIZE to modify process limits,
 3-7
 System console
 ? message, 4-16
 System crash
 See CRASH.COM command procedure
 See Crash dumps
 See System failures
 System directories
 restoring original names
 before upgrading, 3-5
 System disks
 adding an alternate root directory, 2-31
 automatic mounting of, 5-11
 backing up, 5-49, 10-48, 10-55, 10-59
 backing up after installation, 3-13
 backing up for software installations, 3-7

System disks (cont'd)

- booting from an alternate, 4-3
- building with VMSKITBLD, 2-27
- completing a disk created with VMSKITBLD, 2-29
- configuring a system root added with VMSKITBLD, 2-32
- copying system files from, 2-26
- copying system files using VMSKITBLD, 2-29
- disk space needed to run UETP, 17-5
- fragmentation of, 15-27
- installing software on alternate, 3-18
- moving files off to improve system performance, 16-8
- moving page and swap files off, 15-5
- not in volume sets, 8-31
- quotas for, 8-50
- removing and adding optional system files, 5-1
- restoring, 10-57
- saving space by removing optional files, 5-1
- saving space on, 15-5
- test error during UETP, 17-24, 17-25
- testing with UETP, 17-33
- UETP test image, 17-34

System Dump Analyzer utility (SDA)

- analyzing the system dump file, 15-2
 - in system startup, 5-13, 15-4, 15-11
- COPY command, 15-15
- determining cause of system failure, 15-11
- freeing dump information from the page file, 15-16
- saving contents of system dump file, 15-15

System dump files

- See also Crash dumps
- See also SYSDUMP.DMP file
- See also System dump file sizes
- See also System failures
- analyzing, 15-11
- calculating size, 15-6
- comparison of contents of physical and selective dumps, 15-3, 15-10
- copying with BACKUP, 15-4, 15-14
- default location, 15-2
- definition, 15-2
- deleting after creating a new version, 15-28
- displaying the size calculated by AUTOGEN, 15-17, 15-22
- freeing page file, 15-4
- information captured in, 15-2
- installing
 - automatically, 15-2
 - when resized with AUTOGEN, 15-18, 15-22
- insufficient disk space, 15-10
- investigating cause of system failure, 15-11
- lack of, 15-2
- overwriting of, 15-2

System dump files (cont'd)

- protecting with UIC security, 15-4
- requirements, 15-3
 - location, 15-3
 - size, 15-4
- saving contents on reboot, 5-13, 15-15
- saving minimal information in, 15-10
- size
 - See System dump file sizes
- storing selective portions of memory, 15-10
- tasks for managing, 15-1
- use of page file for, 15-2

System dump file sizes

- calculating
 - manually, 15-6
 - with AUTOGEN, 15-2
- changing
 - recommended method, 15-22
 - with AUTOGEN, 15-10, 15-22
 - with SYSGEN, 15-26
- displaying AUTOGEN's calculations, 15-17, 15-22
- minimizing, 15-10
- required, 15-4
 - for page file, 15-4

System failures, 15-2

- See also Crash dumps
- See also System dump files
- determining cause, 15-2, 15-11
- reporting with a Software Performance Report, 15-11
- saving contents of system dump file after, 5-13, 15-15
- writing of system dump file, 15-2

System files

- duplicating using VMSKITBLD, 2-26
- moving off system disk to improve performance, 16-8
- on public volumes, 8-8
- optional

- adding or deleting, 5-2

System Generation utility (SYSGEN)

- and version checking, 5-23
- AUTOCONFIGURE command (VAX)
 - in system startup, 5-7
- changing page, swap, and dump file sizes, 15-22, 15-26
- changing system parameter file with, 14-33
- changing system parameters, 14-33, 18-20
 - See also System parameters, changing
- configuring devices
 - in system startup, 5-7
- CONNECT command (VAX), 7-7
 - in system startup, 5-7
- converting parameters for use with AUTOGEN, 14-5
- CREATE command, 15-18, 15-26

System Generation utility (SYSGEN) (cont'd)

- creating a new system parameter file, 14-34
- creating page, swap, and dump files, 15-18
- INSTALL command, 15-19
 - in SYPAGSWPFILES.COM, 15-21
 - in system startup, 5-7
- installing page, swap, and dump files, 15-19
 - in SYPAGSWPFILES.COM, 15-21
- installing page, swap and dump files
 - in system startup, 5-7
- LOAD command (VAX), 7-7
- managing system parameters, 14-4
- operator log messages, 18-20
- showing system parameters, 14-31

System hangups, 17-21, 17-29

System libraries

- decompressing, 16-6

System management

- centralizing with SYSMAN, 2-10
- clearing counters, 22-25
- creating access control lists (ACLs), 11-9
- deleting device information, 22-17
- deleting node information, 22-20
- disabling MOP downline load service, 22-23
- displaying device information, 22-15
- displaying LAN device configurations, 22-9
- displaying LAN device parameters, 22-10
- displaying node information, 22-18
- displaying OPCOM messages, 22-25
- displaying status and counters, 22-23, 22-24
- enabling MOP downline load service, 22-23
- environment, 2-12 to 2-15
- LANACP device database, 22-14
- LAN Auxiliary Control Program (LANACP) utility, 22-5
- LAN Control Program (LANCP) utility, 22-6
- LANCP command files, 22-9
- LANCP SPAWN function, 22-9
- LAN devices, 22-9
- LAN enhancements, 22-1
- LAN firmware updates, 22-14
- LAN node database management, 22-17
- load trace facility, 22-25
- MOP console carrier, 22-26
- MOP downline load service management, 22-23
- MOP trigger boot, 22-27
- on multiple nodes, 2-13
- running the LANACP utility, 22-6
- running the LANCP utility, 22-7
- setting device information, 22-15
- setting LAN device parameters, 22-12
- setting node information, 22-18
- stopping the LANACP utility, 22-6
- tasks
 - clusterwide management, 20-3
 - establishing node in network, 21-11

System Management utility (SYSMAN)

- accessing disks, 8-50
- adding startup files to a startup database, 5-20, 5-21
- ALF commands, 6-32
- authorization checks in, 2-16
- changing privileges in, 2-17
- changing system parameters
 - active values, 14-28
- command verification in, 2-17
- configuring devices (Alpha)
 - in system startup, 5-7
- converting parameters for use with AUTOGEN, 14-5
- creating command procedures for, 2-18
- deleting startup files, 5-21
- disabling startup files, 5-22
- DISKQUOTA commands, 8-50
- disk quotas with, 8-48
- Disk Quota utility, 8-50
- DO command, 2-18
- enabling remote systems to execute commands, 2-12
- enabling startup files, 5-22
- features of, 2-10
- how commands execute, 2-11
- initialization file, 2-19
- IO AUTOCONFIGURE command (Alpha)
 - in system startup, 5-7
- IO CONNECT command (Alpha), 7-8
 - in system startup, 5-7
- IO LOAD (Alpha), 7-8
- loading licenses with, 3-8
- management environment, 2-12
- managing a VMScluster, 20-16 to 20-23
- managing startup, 5-18
- managing system parameters, 14-4, 14-25
- modifying the system parameter file, 14-28
- PARAMETERS command, 14-25
 - summary, 14-25
- privileges required, 2-11
- profile, 2-16
 - adjusting, 2-16
- showing system parameters, 14-27
- showing the contents of a startup database, 5-20
- showing the name of the target startup database, 5-20
- shutdown, 4-34
- SMISERVER process, 2-12
- specifying the current startup database, 5-20
- STARTUP command, 5-18
- startup logging, 4-15
- startup restrictions, 23-12
- timeout periods, 2-18
- use of passwords, 2-13, 2-16
- using logical names in, 2-14
- using to centralize system management, 2-10

System messages

See also Messages

using when managing a system, 2-5

System parameters

See also Parameter files

ACP cache system, 8-43

active values, 14-3, 14-26

affected by AUTOGEN calculations, 14-10

ALPHAVMSSYS.PAR file (Alpha), 14-3

automatic setting by AUTOGEN, 14-3

booting with default, 4-8

categories by function, 14-2

changing

checking AUTOGEN's settings, 14-10

editing MODPARAMS.DAT file, 14-18

recommended method, 14-4, 14-18

specifying values in MODPARAMS.DAT file, 14-4

with AUTOGEN, 14-18

with conversational boot, 4-3, 4-6, 14-36

with SYSGEN, 14-33

with SYSMAN, 14-28

checking before using VMSINSTAL.COM

command procedure, 3-7

controlling

increasing, 14-19

in MODPARAMS.DAT file, 14-4, 14-18

specifying absolute values, 14-19

specifying maximum values, 14-20

specifying minimum values, 14-20

with ADD_ prefix, 14-19

with MAX_ prefix, 14-20

with MIN_ prefix, 14-20

creating a new parameter file

with SYSGEN, 14-34

current values, 14-3, 14-26

default values, 14-3

definition, 14-1

DUMPBUG, 15-3

DUMPSTYLE, 15-3, 15-10

dynamic, 14-3

effect on other parameters, 14-4

ERLBUFFERPAGES, 15-6

ERRORLOGBUFFERS, 15-6

file extensions, 16-7

GBLPAGES, 16-12

GBLSECTIONS, 16-12

initialization at boot time, 14-36

in memory

See Active system parameters

MODPARAMS.DAT file

See MODPARAMS.DAT file

MULTIPROCESSING, 26-3

MVTIMEOUT, 8-60, 8-62

NPAGEDYN, 26-7

on disk, 14-3

See also Current system parameters

System parameters

on disk (cont'd)

See Current system parameters

in ALPHAVMSSYS.PAR file (Alpha), 14-36

in VAXVMSSYS.PAR file (VAX), 14-36

parameter files

See Parameter files

QUANTUM, 26-8

recommended method for changing, 14-4, 14-5

RMS_EXTEND_SIZE, 16-7

SAVEDUMP, 15-2, 15-3

SCSNODE, 23-10

showing

with conversational boot, 4-3, 4-6, 14-36

with SYSGEN, 14-31

with SYSMAN, 14-27

SMP_CPUS, 26-3, 26-6

STARTUP_P1, 4-14

STARTUP_P2, 4-14

storing your changes for use with AUTOGEN, 14-5

symmetric multiprocessing, 26-3

TAPE_MVTIMEOUT, 8-60, 8-62

tasks for managing, 14-1

TTY_DEFCHAR, 7-10

TTY_DEFCHAR2, 7-10, 7-11

types of, 14-2

dynamic, 14-2

general, 14-2

major, 14-2

UAFALTERNATE, 4-10

user definable, 14-3

VAXVMSSYS.PAR file (VAX), 14-3

vector processing, 26-5

VECTOR_MARGIN, 26-8

VECTOR_PROC, 26-5

VIRTUALPAGECNT, 13-32

when incorrect values prevent the system from booting, 4-8

WSMAX, 13-30, 26-7

System passwords, 11-3

dictionary of, 11-2

System shutdown

See also SHUTDOWN.COM command procedure

adjusting quorum when shutting down a node, 4-29

after software installation, 3-13

allowing batch and print jobs to complete before, 13-55

caution about timing of system halt, 15-2

checking for existence of system files before, 4-29

customizing, 4-32

with SYSHUTDOWN.COM command procedure, 4-27

defining the minimum number of minutes before shutdown, 4-33

System shutdown (cont'd)

- emergency procedure
 - OPCCRAASH, 4-27, 4-36
- emergency procedures
 - console, 4-27
- for an entire VMScluster, 4-29
- notification of, 4-33
- options
 - automatic reboot, 4-29
 - manual reboot, 4-29
 - specifying time interval between DISABLE AUTOSTART/QUEUES and STOP /QUEUES/ON_NODE commands, 13-56
 - time of shutdown, 4-33
- orderly, 4-27
- order of events, 4-31
- procedures for performing, 4-27
- saving AUTOGEN feedback data, 4-29
- SHUTDOWN.COM, 4-27
 - example, 4-30
 - how to use, 4-28
 - when to use, 4-27
- stopping queues before, 13-55
- with SYSMAN, 4-34

System startup

- See also Booting
- See also Site-specific startup command procedure
- See also Startup command procedure
- See also Startup phases
- analyzing a crash dump, 5-13
- assigning systemwide logical names, 5-8
- booting with minimum, 4-14
- CONFIGURE phase, 7-6
- configuring devices, 4-4, 7-6
 - special (Alpha), 5-7
 - special (VAX), 5-7
- controlling when booting, 4-12
- creating systemwide announcements, 5-14
- databases, 5-18
- defining location of systemwide login procedure, 5-17
- definition of logical names, 5-4, 5-5
- description, 4-12
- determining order of phases, 5-18
- displaying startup commands as they execute, 4-14
- enabling autostart, 5-12
- enabling operator console, 5-5
- enabling operator log file, 5-5
- events, 4-4
 - order of, 5-4
 - possibility of future change in order, 5-5
- execution of AUTOCONFIGURE command, 5-4
- execution of login procedures, 5-16

System startup (cont'd)

- execution of site-specific startup command procedures, 5-5
- freeing dump information from page file, 15-4
- in an emergency
 - with default system parameters, 4-8
 - without startup and login procedures, 4-8
 - without the UAF, 4-10
- installing images, 5-4, 16-10
- installing page and swap files, 5-5, 5-6, 15-5, 15-19, 15-21
- limiting the number of interactive users, 5-16
- LMF database, 5-5
- loading of device drivers, 5-4
- loading of licenses, 5-5
- loading of Product Authorization Keys (PAKs), 5-5
- location of files used in, 4-4
- logging with SYSMAN, 4-15
- making remote InfoServer devices available, 23-14
- making remote InfoServer disks available, 5-13
- managing with SYSMAN, 5-18
- messages
 - indicating execution of site-independent startup, 4-5
 - indicating execution of site-specific startup, 4-5
 - indicating lack of installed page file, 5-6
- mounting disk for page and swap files, 5-6
- mounting the queue database disk, 12-6
- performing site-specific operations, 5-10
- phases
 - See Startup phases
- purging the operator log file, 5-14
- saving contents of system dump file, 5-13, 15-15
- setting
 - device characteristics in, 5-12
 - printer device characteristics, 7-13
 - terminal device characteristics, 7-10
- setting up a LAT network, 5-15
- starting InfoServer Client for OpenVMS software, 5-13
- starting of system processes, 5-5
- starting queues, 5-12
- starting SMISERVER process, 5-5
- starting the DECnet network, 5-15
- starting the License Management Facility (LMF), 5-5
- starting the queue manager, 12-3
- startup command procedures, 5-2
- startup of CONFIGURE process, 5-5
- submitting batch jobs, 5-14
- suppressing autoconfiguration, 5-8, 7-8
- tasks, 4-1
- VMS\$PHASES.DAT database, 5-4

- System time
 - See Time
 - resetting after January 1st, 20-18
- System tuning
 - See Tuning
- System users (security category)
 - defining with MAXSYSGROUP parameter, 11-8
 - qualifications for, 11-8
- System version
 - registering images with dependencies on, 5-22
- System volumes
 - definition, 8-8
- Systemwide logical names, 5-8
- SYSTEST account
 - initial modification, 6-9
 - in UAF, 6-7
 - logging into for UETP, 17-2, 17-4
 - privileges required for UETP, 17-23
 - quotas required to run UETP, 17-23
- SYSTEST directory
 - creating for UETP, 17-6
 - function during UETP, 17-4
- SYSTEST_CLIG account
 - in UAF, 6-7
 - reenabling for UETP, 17-10
 - requirements for UETP, 17-10, 17-37
- SYSUAF.DAT files
 - definition, 6-2
 - moving to reduce system disk I/O, 16-8
- SYSUAFALT.DAT file, 4-10
- SYSUAF logical name
 - defining during system startup, 5-8
 - defining to reduce system disk I/O, 16-8

T

- Tailoring a system disk
 - with VMSTAILOR and DECW\$TAILOR, 5-1
- Tailoring utilities (VMSTAILOR and DECW\$TAILOR), 5-1
- Tape cartridge drives
 - preparing for UETP, 17-7
- Tape commands
 - DISMOUNT, 8-44
 - MOUNT, 8-18
- Tape files
 - See also Tape file system
 - accessing at file level, 9-13
 - accessing for read and write operations, 9-14
 - append operation, 9-18
 - closing after opening for read access, 9-17
 - closing after opening for write access, 9-18
 - copying, 9-21
 - definition, 8-7
 - locating for read and write access, 9-15
 - modifying characteristics, 9-8

- Tape files (cont'd)
 - reading, 9-17
 - update operation, 9-18
 - writing to, 9-18
- Tape file system
 - checking
 - continuation volume, 8-41
 - expiration date field, 9-18
 - locating files, 9-15
 - overwriting existing information, 9-17
 - protection on, 8-19
 - writing files to tape volume, 9-18
- Tape mass storage control protocol
 - See TMSCP
- Tapes
 - See also Tape volumes
 - accessing file level, 9-14
 - allocating drives, 8-9, 9-22
 - basic concepts of, 8-6
 - blocks, 8-6
 - bpi, 8-6
 - commands, 8-27
 - copying files from, 9-21
 - deallocating drives, 8-10
 - dismounting, 10-16
 - DOS-11, 9-24
 - enabling write cache, 8-24
 - file names, 9-16
 - file protection
 - See Protection
 - files
 - See Tape files
 - file system, 8-7
 - getting information about, 7-16
 - initializing, 10-13
 - IRG (interrecord gap), 8-6
 - label format, 8-24
 - labeling, 10-21
 - loading on drive, 9-22
 - management
 - tasks, 7-16
 - markers on, 8-6
 - modifying device characteristics, 7-16, 8-42
 - mounting, 10-15
 - MTACP process, 8-6
 - preparing for UETP, 17-2, 17-5, 17-6
 - protection, 8-14
 - reading from, 9-17
 - record blocking, 8-6
 - advantages, 8-7
 - record size
 - specifying, 8-26
 - sequential organization of, 8-6
 - specifying block size, 8-24
 - standard-labeled
 - mounting, 8-24
 - structure of, 8-6

Tapes (cont'd)

- testing with UETP, 17-31, 17-33

- UETP test image, 17-34

- volume label

 - overriding protection, 8-25

- volume protection

 - See Protection

- volumes

 - See also Tape volumes

- volume sets, 8-6

- write cache

 - enabling, 8-24

- writing files to, 9-22

Tape volumes, 9-23

- See also Disk volumes

- See also Tape files

- See also Tapes

- See also Tape volumes

- See also Volumes

- accessibility protection, 8-19

- accessing files on, 9-13, 9-17

- access to, 9-14

- continuation, 8-39

- copying files, 9-21

 - from, 9-21

 - to, 9-20

 - to and from, 9-24

- deallocating, 9-23

- dismounting, 9-23

- file-structured, 8-20

- foreign, 8-20

- header labels, 8-24

- initializing, 9-22

- mounting, 8-21, 8-24, 8-27

- mounting volume sets, 8-36

- mounting with automatic switching disabled, 8-40

- mount verification, 8-58

- overriding UIC, 8-25

- private, 8-9

- reading attributes of header labels, 9-17

- reading files on, 9-17

- searching for files on, 9-5

- specifying record size, 8-26

- standard-labeled, 9-21

 - copying files from, 9-21

- wildcard characters supported, 9-16

- write-enabling, 8-28

- write-locking, 8-37

- write rings, 8-37

- writing files to, 9-22

- writing to files on, 9-18

- TAPE_MVTIMEOUT system parameter, 8-60, 8-62

TCP/IP (Transmission Control Protocol/Internet Protocol)

- applications, 21-13

TDF (time differential factor), 5-34, 5-36

- changing for daylight saving time, 5-35

- definition, 5-34

- how to set, 5-34

- map for determining, 5-36

- tables of, B-1

TELNET, 21-13

Template files

- for site-specific startup, 5-3

Temporary working directory

- specifying alternate working device for, 3-14

Terminal queues, 13-3

Terminals

- console, 2-21

- controlling access through system passwords, 11-3

- determining physical type, 7-12

- disconnecting without terminating a process, 7-11

- documenting owners of, 7-10

- for security alarms, 18-29

- keeping sessions on multiple, 7-11

- LAT, 7-12

 - determining characteristics of, 7-12

 - disconnecting, 7-11

- managing

 - tasks for, 7-9

- operator

 - See Operator terminals

- preparing for UETP, 17-3, 17-5, 17-8, 17-14

- remote, 7-11

- SET TERMINAL/INQUIRE command, 7-12

- setting characteristics, 7-9

 - default values, 7-10

 - in system startup, 5-12, 7-10

- simulating users for UETP, 17-34

- testing with UETP, 17-31, 17-33

- UETP output, 17-33

- UETP test image, 17-34

- virtual

 - See Virtual terminals

Terminal servers, 24-5

- defined, 24-1

- on OpenVMS system, 24-8

Time

- See also TDF (time differential factor)

- modifying system time, 20-17

- resetting after January 1st, 20-18

- updating in a VMScluster, 20-18

Time conversion service, 5-33

Time differential factor

- See TDF

- Time differential factor (TDF)
 - See also Time
- Time formats, 5-42
 - predefined, 5-47
 - specifying, 5-45
- Timeout periods
 - in SYSMAN, 2-18
 - mount verification OPCOM message, 8-60
- Timer queue entry limit, 6-43
- Time zones
 - setting up your system to compensate for, 5-33
- TLZ04 tape drives
 - supported by UETP, 17-7
- TMSCP (tape mass storage control protocol), 20-2
- TN3270, 21-13
- TP_SERVER process
 - stopping permanently, 25-20
 - stopping to dismount a disk, 25-14
- TQELM process limit, 6-43
- Track
 - definition, A-2
- Trailer labels
 - on tape files, 8-7
- Trailer pages, 13-33
 - file, 13-37
 - job, 13-37
- Transaction group
 - example, 25-22
- Transaction logs
 - changing the size of, 25-9
 - checking the size of, 25-8
 - creating, 25-4
 - moving, 25-11
 - planning, 25-3
- Transactions
 - monitoring, 25-6
- Transferring files remotely, 21-14
- Translation modes
 - card reader, 7-18
- Transmission Control Protocol/Internet Protocol
 - See TCP/IP
- Transmission speed
 - setting for terminals, 7-9, 7-10
- Transports, 23-6
 - LASTport, 23-5
- Trigger boot
 - MOP downtime load service, 22-27
- Troubleshooting
 - adding or deleting a device control library module, 13-83
 - autostart queues, 13-81
 - booting problems, 4-14, 4-16
 - general printer problems, 13-78
 - holding jobs, 13-79
 - if a device is not recognized by the system, 7-7
 - jobs that will not execute, 13-79
 - jobs with characteristic mismatch, 13-81
 - Troubleshooting (cont'd)
 - OPCOM problems, 2-22
 - pending jobs, 13-79
 - print jobs with stock mismatch, 13-80
 - problems deleting a queue, form, or characteristic, 13-82
 - queue manager, 12-14
 - queue problems, 13-77
 - stalled output queue, 13-81
 - startup problems, 4-14
 - system dump file for, 15-11
 - system failure, 15-11
 - system hang, 15-16
 - system startup problems, 4-15
 - UETP, 17-21
 - TT2\$M_DISCONNECT characteristic
 - enabling, 7-11
 - relationship to TTY_DEFCHAR2 system parameter, 7-11
 - setting up virtual terminals, 7-11
 - TTDRIVER device driver
 - loading, 7-11
 - TTY_DECCHAR system parameter, 7-10
 - TTY_DEFCHAR2 system parameter, 7-10
 - relationship to TT2Y\$M_DISCONNECT characteristic, 7-11
 - setting up virtual terminals, 7-11
- Tuning
 - See also Performance
 - considering hardware capacity, 16-6
 - definition, 16-4
 - evaluating success, 16-6
 - minimizing with AUTOGEN feedback, 14-10
 - predicting when required, 16-5
 - with AUTOGEN, 14-5
- TYPE command
 - tape, 9-17

U

- UAFALTERNATE logical name, 4-10
- UAFALTERNATE system parameter, 4-10
- UAFs (user authorization files)
 - booting with alternate, 4-10
 - checking quotas for software installation, 3-7
 - description of, 6-2
 - general maintenance, 6-6
 - initial contents, 6-6
 - initial modification, 6-9
 - listing records in, 6-23
 - logical name defining location, 5-8
 - login check, 6-5
 - modifying
 - user record, 6-23
 - network proxy, 6-35
 - records
 - creating multiple default, 6-24
 - resource limits for VAX and Alpha, 6-40

UAFs (user authorization files) (cont'd)

- returning to the default, 4-11
- SYSMAN checks, 2-16
- SYSUAF.DAT, 6-2
- user priorities, 6-31

UETCONT00.DAT file, 17-31

- creation of, 17-31

UETINIDEV.DAT file, 17-31, 17-32

- creation of, 17-31

- format, 17-32

UETININET.DAT file, 17-36

UETINIT00.EXE image, 17-31

UETINIT01.EXE image, 17-21, 17-31

UETLOAD00.DAT file, 17-34

UETNETS00.EXE file, 17-36

UETP\$NODE_ADDRESS logical name, 17-12

UETP (User Environment Test Package)

- aborting execution, 17-15
- compact disc drives supported by, 17-7
- DECnet installation defaults, 17-30
- description, 17-1
- displaying tests as they run, 17-18
- initialization phase, 17-31
- interpreting output, 17-17
- master command procedure, 17-30
- normal completion, 17-15
- optical disk drives supported by, 17-7
- organization, 17-30
- required privileges, 17-23
- required quotas, 17-23
- requirements for small disk systems, 17-11
- running all phases, 17-3
- running individual phase, 17-13
- running multiple passes, 17-14, 17-20
- starting, 17-13
- system resource requirements, 17-2, 17-4
- testing vector processors, 17-12
- TLZ04 tape drive, 17-7
- typical failures reported by, 17-21
- VAX Vector Instruction, 17-12

UETP.COM procedure, 17-30

- terminating, 17-15

UETP.LOG file, 17-15, 17-20, 17-29, 17-35

UETPHAS00.EXE program, 17-30, 17-31

UETUNAS00.EXE UETP test image, 17-19

UFDs (user file directories)

- included in MFD, A-4
- UIC associated with, A-4

UICs (user identification codes), 6-14

- default protection
 - changing, 9-7
- directory protection, 9-11
- disk usage kept in quota file, A-9
- for UETP, 17-6
- identifiers, 11-10
- interpreting, 11-7
- member number, 6-18
- overriding for tape volumes, 8-25

UICs (user identification codes) (cont'd)

protection

- of queues, 13-22
- public volumes, 8-15

- [0,0], 8-49

Update procedures

- See also Mandatory updates

- definition, 3-5

- restrictions, 3-5

Upgrade procedures

- and system version dependent applications, 5-23

- definition, 3-5

- on Alpha systems, 3-2

- POLYCENTER Software Installation utility, 3-2, 3-29

- restrictions, 3-5

- steps in, 3-5

- using with existing OpenVMS, 3-2

Usage count

- DIRECTORY/SIZE command, 8-49

- DISKQUOTA display, 8-49

USE command

- in conversational boot, 4-7

- in SYSGEN, 14-31, 14-33

User accounts

- changing quotas or privileges, 6-23

- deleting, 6-25

- disabling, 6-27

- listing records of, 6-23

- maintaining, 6-24

- modifying, 6-23

- restricting use, 6-27

- setting up, 6-6

User authorization, 6-5

User authorization files

- See UAFs

User disks

- preparing for UETP, 17-2, 17-6

- space requirements for UETP, 17-5

- test error during UETP, 17-24

- testing with UETP, 17-33

- UETP test image, 17-34

User Environment Test Package

- See UETP

User file directories

- See UFDs

User files

- on public volumes, 8-8

- placement, 8-9

User Identification Code

- See UIC

User loads

- defined for UETP DECnet test, 17-36

- defining for the UETP load test, 17-14

- equation used to determine for UETP load test, 17-18, 17-19

- User Mail profile, 6-38
- User mode
 - logical names, 16-15
- User names
 - as identifiers, 11-10
 - entering when logging in, 2-10
- User resources, 6-39
- Users
 - interactive
 - limiting number of, 5-16
 - limiting the number of interactive, 5-16, 16-4
 - protection code categories, 11-8
 - restricting login hours for, 16-4
 - security categories of, 11-8
 - sending messages to with OPCOM, 2-22
 - sending requests to an operator, 2-25
 - validation of, 2-16
- User-specified job retention
 - changing, 13-74
- UTC\$CONFIGURE_TDF.COM command
 - procedure
 - sample session, 5-39
 - showing and setting your time differential factor (TDF), 5-35
- UTC (Coordinated Universal Time), 5-33 to 5-34
 - compensating for differing time zones, 5-34

V

- Validation of users, 2-16
- VAXcluster environments
 - defining the number of nodes for AUTOGEN, 14-21
- VAXcluster systems
 - adjusting quorum after shutting down a node, 4-29
 - shutting down an entire VAXcluster, 4-29
- VAX systems
 - downline loading, 23-2
- VAX Vector Instruction Emulation facility (VVIEF), 17-12
 - definition, 26-5
 - determining presence of, 26-9, 26-10
 - loading, 26-10
 - unloading, 26-10
- VAXVMSSYS.PAR file, 4-2, 14-3
 - initializing parameters at boot time, 14-36
- VCRs (vector count registers), 26-4
- Vector capability
 - determining availability within a system, 26-9
 - placing an ACL on, 26-8
- Vector-capable, 26-4
- Vector consumer
 - determining the identity of, 26-9
 - managing, 26-6
 - marginal, 26-8
 - obtaining information about, 26-8

- Vector context switch
 - obtaining information about, 26-9
- Vector CPU time
 - obtaining information regarding process, 26-9
- Vector-present processors, 26-4
 - adding to system, 26-6
 - identifying, 26-9
 - removing from system, 26-6
 - when unavailable, 26-6
- Vector processing, 26-4 to 26-9
 - configuring system, 26-6
 - definition, 26-4
 - establishing batch queues for, 26-7
 - managing, 26-5
 - obtaining information about, 26-8
 - obtaining number of vector processors, 26-9
 - resource requirements, 26-7
 - system performance, 26-4
 - tasks for managing, 26-5
 - tuning system, 26-7
 - VAX support for, 26-4
- Vector registers, 26-4
- Vectors
 - definition, 26-4
- VECTOR_MARGIN system parameter, 26-8
- VECTOR_PROC system parameter, 26-5
- Verification
 - mount
 - See Mount verification
 - startup, 4-15
 - turning on during system startup, 4-14
- Version dependencies
 - registering images, 5-22
- Version limits
 - setting on files, 8-53
- Version numbers, 9-15
- Virtual devices
 - mounting during system startup, 23-14
- VIRTUALPAGECNT system parameter
 - optimizing batch queues for Sort/Merge utility, 13-32
- Virtual terminal protocol, 21-13
- Virtual terminals
 - connecting, 7-11
 - determining physical terminal type, 7-12
 - enabling, 7-11
 - purpose of, 7-11
 - TT2\$M_DISCONNECT characteristic, 7-11
 - TTY_DEFCHAR2 system parameter, 7-11
- VLRs (vector length registers), 26-4
- VMB.EXE file, 4-18
 - role in boot process, 4-2
- VMRs (vector mask registers), 26-4
- VMS\$LAYERED.DAT file, 5-19
 - function in startup procedure, 5-18
- VMS\$PHASES.DAT file
 - in startup procedure, 5-4, 5-18

VMS\$VMS.DAT file

in startup procedure, 5-18

VMScluster environments

dual-architecture

installing images, 20-21

preparing for UETP, 17-11

security audit log files in, 18-30

test failure during UETP, 17-26

VMScluster systems

See also VAXcluster environments

adjusting quorum after shutting down a node,
4-29

autostart queues in, 13-4

benefits of, 20-1

common parameter files in, 14-22

cross-architecture booting, 22-23

defining in SYSMAN, 2-15

defining location of queue database files, 12-5

device names in, 7-1

disks in, 8-2

dismounting a volume, 8-45

dual-architecture

installing images, 20-20

executing commands in, 2-11

generic batch queues in, 13-7

generic output queues in, 13-12

generic queues in, 13-2

in network environment, 21-6

LAN management enhancements, 22-1

LAN MOP coexisting with DECnet MOP,
22-20

limiting nodes that can run the queue manager,
12-9

local and nonlocal, 2-15

managing with SYSMAN, 20-16 to 20-23

migrating from DECnet MOP to LAN MOP,
22-20

monitoring with SHOW CLUSTER, 20-4

mounting volumes in, 8-21

moving the queue manager to another node,
12-10

node restriction for a startup command
procedure, 5-19

print and batch queues in, 20-2

remote monitoring limitation, 18-43

security, 20-16

setting up for LAN MOP, 22-22

shared resources, 20-1

shutting down an entire VMScluster, 4-29

specifying location of queue database, 12-6

specifying preferred order of queue manager
nodes, 12-9

SYSMAN management environment, 2-12

system management, 20-3

using CLUSTER_CONFIG.COM, 20-2

using CLUSTER_CONFIG to set up LAN MOP,
22-21

VMSINSTAL.COM command procedure

See also Installation procedures

See also Installing software

Alternate System Root option, 3-18
restriction, 3-18

Alternate Working Device option, 3-14

answer file, 3-14

Autoanswer option, 3-14

BACKUP qualifiers, 3-16

command line syntax, 3-9

completing installation, 3-13

correcting problems, 3-9

creating new answer file, 3-14

destination parameter, 3-12

File Log option, 3-17

Get Save Set option, 3-15, 3-16

getting help in, 3-9

installing layered products, 3-6

options, 3-13

option list parameter, 3-12

specifying, 3-12

table of, 3-11

preparing to use, 3-6

product list parameter, 3-9

product save-set format, 3-16

Release Notes option, 3-17

saving answers, 3-14

source parameter, 3-11

starting, 3-9

system failure

conditions for, 3-13

system shutdown following, 3-13

temporary working directory

specifying location of, 3-14

VMSKITBLD.COM command procedure

ADD option, 2-31

BUILD option, 2-27

compared to CLUSTER_CONFIG.COM
command procedure, 2-31

completing a system disk created with the
BUILD option, 2-29

configuring a system root added with, 2-32
copying system files from the system disk,
2-26

COPY option, 2-29

options, 2-26

reliance on .TEMPLATE version of site-specific
command procedures, 5-3

VMSMAIL logical name

defining to reduce system disk I/O, 16-8

VMSMAIL_PROFILE.DATA file, 6-38

moving to reduce system disk I/O, 16-8

VMSMAIL_PROFILE logical name

defining during system startup, 5-8

VMSTAILOR

See Tailoring utilities

- Volatile databases
 - network, 21-11
 - VOLPRO privilege, 8-14
 - VOLSET.SYS file
 - See Volume set list file
 - Volume identifier field, 8-39
 - Volume integrity, 8-58
 - Volume labels
 - assigning to devices, 5-11
 - changing, 3-33
 - definition, 10-56
 - format, 8-24
 - specifying for magnetic tape, 10-14
 - used with BACKUP command, 10-56
 - Volume protection
 - See Protection, volume
 - Volume Protection Override (VOLPRO)
 - See VOLPRO privilege
 - Volumes
 - See also Disk volumes
 - See also Private volumes
 - See also Public volumes
 - See also Tape volumes
 - See also Volume sets
 - adding to an existing disk set, 8-34
 - advantages of using separate, 8-30
 - availability
 - OPCOM message, 8-59
 - binding into volume set, 8-22
 - canceling mount verification, 8-62
 - continuation, 8-41
 - controlling cache size, 8-58
 - disabling mount messages, 8-23
 - disk
 - See Disk volumes
 - dismounting, 8-43, 10-16
 - restriction, 16-18
 - Files-11, 8-3
 - foreign
 - copying files to and from, 9-24
 - mounting, 8-21, 9-13
 - group, 8-8
 - initializing, 8-10, 10-13
 - mounting, 8-28, 10-15
 - See also MOUNT command
 - continuation tape, 8-40
 - if device is unavailable, 8-27
 - operator functions, 8-18
 - public, 8-20
 - steps, 8-27
 - mount verification aborted
 - OPCOM message, 8-63
 - mount verification timeout, 8-60
 - OPCOM message, 8-60
 - operator-assisted mount, 8-28
 - private, 8-9
 - Volumes (cont'd)
 - protection, 8-14
 - public
 - See Public volumes
 - rebuild
 - determining if needed, 7-4
 - rebuilding, 8-50
 - recovering from errors, 8-60
 - sharing, 8-23
 - substituting, 8-27
 - system
 - See System volumes
 - tape
 - See Tape volumes
 - Volume security profile
 - reserved file, A-9
 - SECURITY.SYS, A-9
 - Volume set list file
 - definition, A-9
 - reserved file, A-9
 - used by Analyze/Disk_Structure utility, A-9
 - VOLSET.SYS, A-9
 - Volume sets
 - backing up, 10-40
 - CD-ROM
 - partially mounted ISO 9660, 8-35
 - characteristics, 8-30
 - creating, 8-22, 8-30
 - definition, 8-2, 8-30
 - disk, 8-31, 8-34
 - accessing, 8-31, 8-32
 - adding to, 8-34
 - adding volumes, 8-34
 - assigning name to, 8-31
 - creating, 8-31
 - from existing volumes, 8-33
 - from new volumes, 8-32
 - creating files on, 8-32
 - directory structure, 8-31
 - index file, 8-32
 - mounting, 8-31, 8-32, 8-33
 - naming, 8-31
 - information stored in VOLSET.SYS, A-9
 - mounting, 8-31
 - See also MOUNT command
 - privileges to access, 8-32
 - processing continuation volumes, 8-37
 - restoring, 10-47
 - restriction for system disk, 8-31
 - root volume, 8-31
 - tape, 8-6, 8-36, 8-40
 - automatic volume switching, 8-39
 - continuation volumes, 8-38
 - creating, 8-37
 - mounting, 8-36
 - mounting continuation volumes, 8-40

- Volume sets
 - tape (cont'd)
 - mounting with automatic switching disabled, 8-40
- Volume shadowing, 10-40, 10-47
 - minimerge (assisted merge), 10-42
 - mounting disks in host-based shadow set, 10-42
 - RAID, 10-40
- Volume structure
 - ISO 9660, 8-3
- Volume switching
 - automatic, 8-39
- VTAO device
 - connecting, 7-11
- VVIEF
 - See VAX Vector Instruction Emulation facility
- VVIEF\$DEINSTAL.COM command procedure, 26-10
- VVIEF\$INSTAL.COM command procedure, 26-5, 26-10
- VVIEF (VAX Vector Instruction Emulation facility)
 - See VAX Vector Instruction Emulation Facility

W

- Waiting job status, 13-69
- WANs (wide area networks)
 - connections, 21-8
 - multipoint, 21-9
 - point-to-point, 21-8
- Welcome messages
 - defining, 5-14
 - displaying, 5-15
 - login, 2-10
- Wide area networks (WANs)
 - See WANs
- Wildcard characters
 - asterisk (*)
 - using with tape volumes, 9-16
 - in file names, 9-15
 - with OpenVMS extended names, 9-16
 - with standard names, 9-16
 - with tape volumes, 9-15
- Working directory
 - temporary
 - in VMSINSTAL.COM command procedure, 3-14
- Working sets
 - adjusting for vectorized applications, 26-7
 - default size, 6-43
 - extent, 6-44
 - limits and quotas
 - choosing a value for batch queues, 13-30, 13-31
 - specifying for batch queues, 13-21, 13-29
 - specifying for output queues, 13-21

- Working sets (cont'd)
 - quota, 6-44
- Work load
 - adjusting system parameters for, 16-5
 - distributing, 16-4
 - designing efficient applications, 16-4
 - restricting login hours, 16-4
 - importance of knowing, 16-2
 - monitoring, 16-2
 - tools used for, 16-2
 - strategies for managing, 16-3
- Workstations
 - backing up, 10-36
 - managing queues on, 13-2
 - OPCOM behavior on, 2-21
 - OPCOM startup on, 2-21
 - printer queue configuration, 13-8
 - setting up media on, 8-11
 - starting SMISERVER process, 2-12
 - use of Snapshot facility with, 4-21
- World users (security category), 11-8
- Writable images, 16-11
- Write access
 - See Access types, write
 - for disk directory files, 9-11
 - for disk files, 9-6
 - granting through protection codes, 11-8
- Writeboot utility (WRITEBOOT), 4-17
 - error messages, 4-19
- Write cache
 - enabling for tape device, 8-24
- Write-enabling a tape, 8-28
- Write-locked devices
 - mount verification, 8-61
- Write-locking
 - disk volumes, 8-60
 - tape volumes, 8-37
- Write operation
 - See Access types, write
- Write rings
 - on tape volumes, 8-37
- WSDEFAULT process limit
 - choosing a value for batch queues, 13-30
 - specifying a value for batch queues, 13-21, 13-29
 - specifying a value for output queues, 13-21
 - specifying values, 13-21 to 13-30
- WSDEF process limit, 6-43
- WSEXTENT process limit, 6-44
 - choosing a value for batch queues, 13-30
 - for efficient sorting, 13-31
 - specifying a value for batch queues, 13-22, 13-29
 - specifying a value for output queues, 13-22
 - value for efficient backups, 10-10

WSMAX system parameter, 13-30, 26-7

WSQUO process limit, 6-44

WSQUOTA process limit

- choosing a value for batch queues, 13-30

- specifying a value for batch queues, 13-22, 13-29

- specifying a value for output queues, 13-22

- value for efficient backups, 10-10

X

XAR keyword

- with MOUNT/PROTECTION command, 8-23

XARs (extended attribute records)

- protection fields, 8-17

- mount option, 8-17

X terminal client, 23-4

1. The first part of the report
describes the general situation
of the country and the
main problems which
are facing it. It also
mentions the main
achievements of the
government in the
last few years.

2. The second part of the report
deals with the economic
situation of the country.
It mentions the main
problems which are
facing the economy
and the measures
which are being
taken to solve them.

